Optimal Efficiency-Envy Trade-Off via Optimal Transport

Anonymous Author(s)
Affiliation
Address
email

Abstract

We consider the problem of allocating a distribution of items to n receivers where each receiver has to be allocated a fixed, prespecified fraction of all items, while ensuring that each receiver does not experience too much envy. We show that this problem can be formulated as a variant of the semi-discrete optimal transport (OT) problem, whose solution structure in this case has a concise representation and a simple geometric interpretation. Unlike existing literature that treats envy-freeness as a hard constraint, our formulation allows us to *optimally* trade off efficiency and envy continuously. Additionally, we study the statistical properties of the space of our OT based allocation policies by showing a polynomial bound on the number of samples needed to approximate the optimal solution from samples. Our approach is suitable for large-scale fair allocation problems such as the blood donation matching problem, and we show numerically that it performs well on a prior realistic data simulator.

1 Introduction

In this work, we focus on the problem of finding an allocation policy that divides a pool of items, represented by a distribution $\mathcal{D}$, to n receivers, under the constraint that each receiver i must be allocated a prespecified fraction p_i^* of the items, where $p_i^* \in (0,1), 1^\top p^* = 1$ is an input to the problem that characterizes the priority of each receiver. We refer to $\{p_i^*\}_{i=1}^n$ as the target matching distribution. In addition to this matching distribution constraint, we also require that the receiver's envy, which we will define formally later, be bounded. Unlike the existing resource allocation literature, where envy-free is either treated as a hard constraint or not considered at all, we allow the central planner to specify the level of envy that is tolerated, and find the most efficient allocation given the amount of envy budget. This allows us to reduce the envy significantly without paying the full price of fairness with respect to efficiency.

To concretely motivate our model, let us consider the blood donor matching problem that was first studied in [18]. The Meta platform has a tool called *Facebook Blood Donations*, where users who opt in to receive notifications are notified about blood donation opportunities near them. Depending on the user's and the blood bank's specific characteristics (e.g., age, occupation for the user, and locations, hours for the blood bank), notifications about different donation opportunities have different probabilities of resulting in an actual blood donation. The platform would like to send each user the most relevant notifications (to maximize the total number of potential blood donations), while maintaining certain fairness criteria for all the blood banks that participate in this program. Although the platform can theoretically send each user multiple notifications about multiple blood banks, for user experience and other practical reasons, this is not done, at least in the model introduced in [18]. Therefore in this problem, users' attention is the scarce resource that the platforms needs to allocate to different blood banks. The most natural type of fairness criteria in this setting is perhaps the number

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

of users that received notifications about each of the blood banks. For example, it is not desirable to match zero users to a particular, potentially inconveniently-located, blood bank, even if matching zero users to this blood bank results in more blood donations in expectation.

Another fairness desideratum commonly studied in the literature is called *envy-freeness*. We say that receiver A envies receiver B if A values B's allocation more than her own. Intuitively, an allocation that is envy-free–where no agent envies another agent–is perceived to be fair.

This paper considers both of the two fairness criteria mentioned above: we study a setting where the goal is to maximize social welfare under a matching distribution constraint, while ensuring that each receiver has bounded envy. We make the following contributions in regards to this problem:

1. We formulate it as a constrained version of a semi-discrete optimal transport problem and show that the optimal allocation policy has a concise representation and a simple geometric structure. This is particularly attractive for large-scale allocation problems, due to the fast computation of a match given an item. This insight also shines new light on the question of when envy arises, and when the welfare price on envy-freeness is large.
2. We propose an efficient stochastic optimization algorithm for this problem and show that it has a provable convergence rate of $O(1/\sqrt{T})$.
3. We investigate the statistical properties of the space of our optimal transport based allocation policies by showing a Probably Approximately Correct (PAC)-like sample complexity bound for approximating the optimal solution given finite samples.

In Section 3 we formally define the problem we are interested in. In Section 4, we show that this problem can be formulated as a semi-discrete optimal transport problem, whose solution has a simple structure with a nice geometric interpretation. Section 5 develops a stochastic optimization algorithm. In Section 6, we show that an ϵ-approximate solution can be found with high probability given $\tilde{O}(\frac{n}{\epsilon^2})$ samples, where n is the number of receivers. Finally, in Section 7 we demonstrate the effectiveness of our approach using both artificial and a semi-real data.

2 Literature Review

Blood donation matching. McElfresh et al. [18] introduced this problem and modeled it as an online matching problem, where the matching quality between an user and a blood bank is assumed to be known to the platform. The model formulation there is complex, as it takes into account the fact that not every blood bank is in need of blood every day, a dynamic which we do not consider here. Furthermore, their matching policy is rather cumbersome, requiring a separate parameter for each (donor, receiver) pair. Compared to their paper, we are able to provide better structural insights to the problem by utilizing a simpler model that still captures the most salient part of the problem.

Online Resource Allocation. Another strand of work that our paper is closely related to is that of online resource allocation, especially those with i.i.d. or random permutation input models [14, 11, 12]. Agrawal, Wang, and Ye [2] studied the setting with linear objective and gave competitive ratio bounds. Then, Agrawal and Devanur [1] generalized the results to concave objectives and convex constraints. Later, Devanur et al. [12] improved the approximation ratio bounds and relaxed the input assumptions on the budgets. Balseiro, Lu, and Mirrokni [5] show that online mirror descent on the dual multpliers does well under both i.i.d. adversarial, and certain non-stationary input settings. However, none of theses papers study the envy-free criterion. Recently, Balseiro, Lu, and Mirrokni [4] studied an online resource allocation problem with fairness regularization. Although the authors did not explicitly study envy regularization, their regularization framework can be modified to accommodate envy regularization. However, like all the other papers mentioned in this paragraph, the offline solution is used as the benchmark to measure regret, but no explicit solution is given to the offline problem. Our analysis focuses on the offline problem, and draws an explicit connection to optimal transport, which allowed us to provide a novel PAC-like analysis on the sample complexity of the problem. In another recent paper, Sinclair, Banerjee, and Yu [23] studied the trade-off between minimizing envy and minimizing waste, which refers to un-allocated resources. Despite close similarity between our titles, their offline benchmark is the standard Eisenberg-Gale program, which is envy-free, but does not address the welfare cost of achieving envy-freenes.

Fair Division. Offline resource allocation, or commonly known as fair division, is also a popular research area. A large body of these papers are formulated as a cake-cutting problem ([21, 9, 20]) where the resources are modeled as an interval and the agents' valuations are represented as functions on this interval. However most existing results are based on relatively simple (e.g., piece wise uniform, or piece wise linear) valuation functions, where as our problem can be thought of as a cake cutting problem with arbitrarily complex valuation functions. Other fair division literature that studies envy-freeness as a fairness criterea usually treat it as a hard constraint, and tries to maximize social welfare subject to that constraint ([10, 25]). Others have also characterized the worst-case loss in social welfare due to the requirement of envy-freeness [8], and other fairness notions [7]. Unlike these papers that focus on achieving zero envy, we treat the allowable envy as a parameter, and find the most efficient solutions subject to the desired amount of envy.

2.1 Background on Optimal Transport

Since our work draws an explicit connection to optimal transport (OT), we provide a summary of key OT results here. Let α, β be two probability measures on the metric spaces $\mathcal{X}, \mathcal{Y}$ respectively. We define $\Pi(\alpha, \beta)$ as the set of joint probability measures on $\mathcal{X} \times \mathcal{Y}$ with marginals α and β. The *Kantorovich formulation* of the optimal transport problem [16] can be written as

$$L(\alpha, \beta) := \min_{\pi \in \Pi(\alpha,\beta)} \int_{\mathcal{X}\times\mathcal{Y}} c(x, y) d\pi(x, y) \tag{1}$$

where $c(x, y)$ is the cost associated with "moving" x to y. This is called a transportation problem because the conditional probability $\pi(y|x)$ specifies a transportation plan for moving probability mass from $\mathcal{X}$ to $\mathcal{Y}$. If β is a discrete measure, i.e. $\mathcal{Y}$ is finite, then it is known [3] that the dual to (1) can be written as (here we abuse the notation β to also represent the vector of probability masses, where β_i is the probability mass on point y_i):

$$\max_{g \in \mathbb{R}^n} \mathcal{E}(g) := \sum_{i \in [n]} \int_{\mathbb{L}_{y_i}(g)} c(x, y_i) - g_i \, d\alpha(x) + g^\top \beta \tag{2}$$

where $n = |\mathcal{Y}|$, and $\mathbb{L}_{y_i}$ is what is sometimes referred to as the *Laguerre cell*:

$$\mathbb{L}_{y_i}(g) = \{x \in \mathcal{X} : \forall i \neq j, c(x, y_i) - g_i \leq c(x, y_j) - g_j\} \tag{3}$$

Proposition 1 (Proposition 2.1 [3]). *If α is a continuous measure, and β a discrete measure, then $L(\alpha, \beta) = \max_g \mathcal{E}(g)$, and the optimal solution π of (1) is given by the partition $\{\mathbb{L}_{y_i}(g^*), i \in [n]\}$, i.e. $d\pi(x, y_i) = d\alpha(x)$ if $x \in \mathbb{L}_{y_i}(g^*)$, 0 otherwise.*

3 Problem Formulation

There is a set of n receivers $\mathcal{Y}$. There is a "pool" of items, represented by a distribution α over $\mathcal{X} \subseteq [0, \bar{x}]^n$. Each random draw from this distribution $X \sim \alpha$ is a vector representing the n receivers' valuations of this item. The goal is to maximize the expected matched utilities of the recipients, while maintaining the constraint that the receiver y_i is matched p_i^* fraction of the times in expectation. Here $\{p_i^*\}_{i=1}^n$ is called the target matching distribution which intuitively represents receivers' importance. A matching policy π takes a valuation vector and maps it (potentially with randomness) to one of the n receivers. Let $\pi(y|x)$ denote the probability of matching the item to y given valuation vector x. The basic problem formulation is to solve the following optimization problem:

$$\begin{aligned} &\max_{\pi} \mathbb{E}_{X\sim\alpha}\left[\sum_{i=1}^{n} X_i \pi(y_i|X)]\right] \\ &s.t. \quad \mathbb{P}\left[\pi(y_i|X)\right] = p_i^* \quad \forall i \in [n] \end{aligned} \tag{4}$$

We assume that $p_i^* > 0$ for all i. This is WLOG, because we can always pretend that a receiver does not exist, if the target fraction for that receiver is 0. An example of such problem can be see in Figure 1, where α is a distribution over the unit square, and the goal is to partition the square into blue and orange regions (given to A and B respectively) such that each region covers the desired p_A^*, p_B^* probability mass. Note that the orange and blue regions are allowed to over lap (probabilitistic

Figure 1: Left: An illustration of what Laguerre cells look like when $n = 2$. Consider any distribution on the support $[0, 1]^2$. The optimal division of the space is to move the diagonal line up or down until the probability mass contained in the orange region is equal to p^*. Right: A pictorial proof of the optimality of such partition. Suppose one can find an ϵ mass above this line that is matched to A, and an ϵ mass below the line that is matched to B, then switching the assignments of these two regions increases the matched weights.

partition), and that the boundary does not have to be linear as illustrated in the figure. As we will show later in Section 4, despite the large design space permitted by the formulation in (4), we can in fact focus on a much smaller design space.

In resource allocation problems, it is often the case that we care not just about efficiency (maximizing social welfare, or the sum of all receivers' utilities), but also other fairness criteria. One of the most commonly studied fairness criteria is envy-freeness. Agent y_i envies another agent y_j if agent y_i values the allocation given to y_j more (after adjusting for their priority weights). We can formally define agent y_i's envy as

$$Envy(y_i) = \max_j \mathbb{E}_\alpha \left[X_i \pi(y_j|X) \frac{p_i^*}{p_j^*} - X_i \pi(y_i|X) \right] \tag{5}$$

Instead of the vanilla formulation in (4), we consider the following more general formulation:

$$\begin{aligned} &\max_\pi \mathbb{E}_{X\sim\alpha} \left[\sum_{i=1}^n X_i \pi(y_i|X)] \right] \\ &s.t. \quad \mathbb{P}\left[\pi(y_i|X)\right] = p_i^* \quad \forall i \in [n] \\ &\qquad Envy(y_i) \leq \lambda_i \quad \forall i \end{aligned} \tag{6}$$

In existing fair resource allocation literature, people focus on finding allocations such that $Envy(y_i)$ is at most 0 for every y_i. This can be a very restrictive constraint, often satisfied at the cost of reducing efficiency by a significant amount (This reduction is sometimes referred to as the Cost-of-Fairness). We take a different approach, and allow the central planner to set non-negative constraints on envy.

4 Optimal Solution Structure

The space of feasible solutions for (6) is large, which makes the problem difficult to optimize directly. However we can use the tools from OT to reduce the search space to something with much more structure. The key observation is that (6) can be formulated as variation of the semi-discrete optimal transport problem given in Equation (1).

Let's first consider the simpler case in (4) where there are no envy constraints. In this case, the problem can be stated in the form of (1) as follows: the cost function is the negative utility of the matched receiver $c(x, y_i) = -x_i$, the β measure is the discrete measure $\sum_{i=1}^n p_i^* \delta_{y_i}$, and the matching policy $\pi(y|x)$ in (6) is exactly the conditional probability of the joint distribution in (1). From Theorem 1 it follows that the optimal matching policy is represented by Laguerre cells given in (3): x is matched to y_i if $i = \arg\min_k -x_k - g_k$. Note that the dual variables $g \in \mathbb{R}^n$ serve as an "adjustment" over the agents' reported utilities, and the resulting matching policy is simply a greedy policy over this adjusted valuation vector.

Geometrically, each Laguerre cell is simply the intersection of half-spaces: $\mathbb{L}_i(g) = \cap_k \{x : x_i + g_i \geq x_k + g_k\}$. To visualize this better, consider the simple setting with two receivers A, B where their

valuations for an item is a joint distribution supported on $[0,1]^2$. Suppose we want to match p^* fraction of the items to receiver B. Figure 1 gives a proof-by-picture that the optimal strategy is to divide the space up with a slope-1 diagonal line such that the probability mass lying above the line is equal to p^*. This geometric interpretation of the matching policy plays a crucial role in getting us a sample complexity bound later in Section 6.

With this geometric interpretation of the solution space in mind, let us consider the more general case with envy constraints as formulated in (6). The envy constraints can be added to the OT problem in (1) like so:

$$L(\alpha,\beta,\lambda) = \min_{\pi\in\Pi(\alpha,\beta)} \int_{\mathcal{X}\times\mathcal{Y}} c(x,y)d\pi(x,y) \tag{7}$$
$$s.t. \int_{\mathcal{X}} c(x,y_j)d\pi(x,y_j) - \int_{\mathcal{X}} c(x,y_j)d\pi(x,y_k)\frac{\beta_j}{\beta_k} \le \lambda_j \qquad \forall (j,k)\in[n]^2, j\neq k$$

Although envy constraints make the solution space more complicated, we show that it retains the geometric structure of being the intersection of half-spaces. The dual of (7) can be derived using Fenchel-Rockafellar's theorem:

$$\max_{g\in\mathbb{R}^n,\gamma\in\mathbb{R}_+^{n^2-n}} \mathcal{E}(g,\gamma) \coloneqq \sum_{j\in[n]} \int_{\mathbb{L}_{y_j}(g,\gamma)} \bar{g}_{\gamma,c}(x,y_j)d\alpha(x) + g^\top\beta - \sum_{j,k,j\neq k}\gamma_{jk}\lambda_j \tag{8}$$

where

$$\bar{g}_{\gamma,c}(x,y_j) \coloneqq \left(1+\sum_{k\neq j}\gamma_{jk}\right)c(x,y_j) - \sum_{k\neq j}\gamma_{kj}c(x,y_k)\frac{\beta_k}{\beta_j} - g_j, \tag{9}$$

$$\mathbb{L}_y(g,\gamma) \coloneqq \left\{x\in\mathcal{X} : y = \arg\min_{y'\in\mathcal{Y}} \bar{g}_{\gamma,c}(x,y')\right\}. \tag{10}$$

Theorem 1. *If α is a continuous measure, and β a discrete measure, then $L(\alpha,\beta,\lambda) = \max_{g,\gamma}\mathcal{E}(g,\gamma)$, and the optimal solution π of (7) is given by the partition $\{\mathbb{L}_{y_i}(g^*,\gamma^*), i\in[n]\}$: $d\pi(x,y_i) = d\alpha(x)$ if $x\in\mathbb{L}_{y_i}(g^*,\gamma^*)$, 0 otherwise.*

Note that when $c(x,y_i) = -x_i$, $\bar{g}_{\gamma,c}(x,y_j)$ is linear in x, which means that the new Laguerre cells $\mathbb{L}_y(g,y)$ given in Equation (10) are still intersections of half spaces (some examples are given later in Figure 2). Furthermore, the allocation policy can be interpreted as a greedy policy based on the adjusted utility given by (9), which contains additional interaction terms that take envy into account.

5 Stochastic Optimization

Theorem 1 shows that the optimal solution to our allocation problem in (6) is represented by the optimal dual solution from (8). To solve this optimization problem, first note that the objective function $\mathcal{E}(g,\gamma)$ is concave. To see this, we can rewrite the objective function as follows:

$$\mathcal{E}(g,\gamma) = \int_{\mathcal{X}} \min_{i\in[n]} \bar{g}_{\gamma,c}(x,y_i)d\alpha(x) + g^\top\beta - \sum_{j,k,j\neq k}\gamma_{jk}\lambda_j \tag{11}$$

Since $\bar{g}_{\gamma,c}(x,y_j)$ is linear in g and γ and taking a minimum preserves concavity, the objective function is concave. Therefore, the dual problem is a constrained convex optimization problem. The gradient of $\mathcal{E}(g,\gamma)$ can be computed as follows:

$$\nabla_g\mathcal{E}(g,\gamma)_j = -\int_{\mathbb{L}_{y_j}(g,\gamma)} d\alpha(x) + \beta_j \tag{12}$$

$$\nabla_\gamma\mathcal{E}(g,\gamma)_{jk} = \int_{\mathbb{L}_{y_j}(g,\gamma)} c(x,y_j)d\alpha(x) - \int_{\mathbb{L}_{y_k}(g,\gamma)} c(x,y_j)\frac{\beta_j}{\beta_k}\,d\alpha(x) - \lambda_j \tag{13}$$

Calculating this gradient is hard, as it involves integration over an arbitrary measure α. However, an unbiased, stochastic version of the gradient can be easily obtained from a single sample $x\sim\alpha$:

$$\hat{\nabla}_g\mathcal{E}(g,\gamma)_j = -\mathbb{1}[x\in\mathbb{L}_{y_j}(g,\gamma)] + \beta_j \tag{14}$$

$$\hat{\nabla}_\gamma\mathcal{E}(g,\gamma)_{jk} = c(x,y_j)\mathbb{1}[x\in\mathbb{L}_{y_j}(g,\gamma)] - c(x,y_j)\frac{\beta_j}{\beta_k}\mathbb{1}[x\in\mathbb{L}_{y_k}(g,\gamma)] - \lambda_j \tag{15}$$

Algorithm 1: Projected SGD for Envy Constrained Optimal Transport

Input: Distribution α, target matching distribution p^*, timesteps T.

Initialize $g_0 = 0, \gamma_0 = 0, \eta = \frac{1}{\sqrt{T}}$.

for $t \leftarrow 0, 1, 2, \ldots, T$ **do**

  Sample $x_t \sim \alpha$

  $g_{t+1} \leftarrow g_t + \eta \hat{\nabla}_g \mathcal{E}(g, \gamma)$

  $\gamma_{t+1} \leftarrow \left(\gamma_t + \eta \hat{\nabla}_\gamma \mathcal{E}(g, \gamma)\right)^+$

return $\sum_{t=1}^T g_t/T, \sum_{t=1}^T \gamma_t/T$

The details of the algorithm is given in Algorithm 1. Standard projected SGD analysis (see for example [15]) tells us that Algorithm 1 converges at the rate $\mathcal{E}(g^*, \gamma^*) - \mathcal{E}\left(\mathbb{E}[g_T], \mathbb{E}[\gamma_T]\right) \leq O\left(\frac{1}{\sqrt{T}}\right)$.

6 Learning from Samples

So far we have considered the setting where the true underlying distribution is known, and assumed that we can freely draw independent samples from that distribution. In many settings, we only have access to α in the form of finite number of i.i.d. samples. In this section, we focus only on the assignment cost function $c(x, y_i) = -x_i$, which models our original resource allocation problem proposed in Section 3. The goal of this section is to establish a sample complexity bound for solving the dual problem (8). Let $S = \{X^1, X^2, \ldots, X^m\}$ be m independent samples from α. The empirical version of the dual objective (11) is:

$$\mathcal{E}_S(g, \gamma) = \frac{1}{m} \sum_{t=1}^{m} \min_{i \in [n]} \bar{g}_{\gamma,c}(X^t, y_i) + g^\top \beta - \sum_{j,k,j \neq k} \gamma_{jk} \lambda_j \tag{16}$$

Let $\hat{g}_S, \hat{\gamma}_S$ be the empirical maximizer given the set of samples S: $(\hat{g}_S, \hat{\gamma}_S) := \arg\max \mathcal{E}_S(g, \gamma)$, and g^*, γ^* be the population maximizer $(g^*, \gamma^*) = \arg\max \mathcal{E}(g, \gamma)$. We are interested in bounding the number of samples needed so that $\mathcal{E}(g^*, \gamma^*) - \mathcal{E}(\hat{g}_S, \hat{\gamma}_S)$ is small with high probability. Let's introduce some new notations to facilitate our later discussions. Define the following hypothesis class for each i:

$$F_i = \left\{ x \mapsto \bar{g}_{\gamma,c}(x, y_i) + g^\top \beta - \sum_{j,k,j \neq k} \gamma_{jk} \lambda_j : g \in \mathbb{R}^n, \gamma \in \mathbb{R}_+^{n(n-1)} \right\}. \tag{17}$$

as well as the overall hypothesis class:

$$F = \left\{ x \mapsto \min_{i \in [n]} \bar{g}_{\gamma,c}(x, y_i) + g^\top \beta - \sum_{j,k,j \neq k} \gamma_{jk} \lambda_j : g \in \mathbb{R}^n, \gamma \in \mathbb{R}_+^{n(n-1)} \right\}. \tag{18}$$

Plugging $c(x, y_i) = -x_i$ into the definition of $\bar{g}_{\gamma,c}$, we see that for a given g, and γ, the corresponding hypothesis $f_i \in F_i$ can be written as $f_i(x) = w^\top x + b$, where

$$w_j = \begin{cases} -(1 + \sum_{k \neq i} \gamma_{ik}), & \text{if } j = i \\ \gamma_{ji} \frac{\beta_j}{\beta_i}, & \text{if } j \neq i \end{cases}, \tag{19}$$

and

$$b = -g_i + g^\top \beta - \sum_{j,k,j \neq k} \gamma_{jk} \lambda_j. \tag{20}$$

It also follows that

$$F \subseteq F_{min} := \{x \mapsto \min_i f_i(x) : f_i \in F_i\}.$$

This interpretation of the original hypothesis class as the minimum over n affine hypothesis classes is the key observation to prove the sample complexity bound. We prove our main result under the following boundedness assumption:

Assumption 1. *The hypothesis $f(x) = \min_i f_i(x) = \min_i {w^i}^\top x + b^i$ corresponding to the optimal dual solution g^*, γ^* satisfies $\|w^i\|_1 \vee |b^i| \leq R$ for some $R > 0$. In particular, these assumptions imply that F_i and F are uniformly bounded by $R\bar{x} + R$.*

From (19) and (20) we can see that this is essentially a bound on the optimal dual variables g^*, γ^*, and a bound on the ratio β_j/β_i, both of which are determined by the input distributions α, β, and do not depend on the number of samples. In other words, R is a problem dependent constant.

Theorem 2. *Under Assumption 1, for a given sample size m, with probability $1-\delta$, $\mathcal{E}(g^*, \gamma^*) - \mathcal{E}(\hat{g}_S, \hat{\gamma}_S) < O\left(\sqrt{\frac{(\log m)^3 + \log(1/\delta)}{m}}\right)$.*

Proof. We prove the result via uniform convergence:

$$\begin{aligned}
&\mathcal{E}(g^*, \gamma^*) - \mathcal{E}(\hat{g}_S, \hat{\gamma}_S) \\
=&\mathcal{E}(g^*, \gamma^*) - \mathcal{E}_S(\hat{g}_S, \hat{\gamma}_S) + \mathcal{E}_S(\hat{g}_S, \hat{\gamma}_S) - \mathcal{E}(\hat{g}_S, \hat{\gamma}_S) \\
\leq&\mathcal{E}(g^*, \gamma^*) - \mathcal{E}_S(g^*, \gamma^*) + \mathcal{E}_S(\hat{g}_S, \hat{\gamma}_S) - \mathcal{E}(\hat{g}_S, \hat{\gamma}_S) \\
\leq& \sup_{g,\gamma} \left(\mathcal{E}(g,\gamma) - \mathcal{E}_S(g,\gamma)\right) + \sup_{g,\gamma} \left(\mathcal{E}_S(g,\gamma) - \mathcal{E}(g,\gamma)\right) \\
\leq& 2 \sup_{g,\gamma} |\mathcal{E}(g,\gamma) - \mathcal{E}_S(g,\gamma)|
\end{aligned} \tag{21}$$

Clearly, it suffices to show that $\mathcal{E}_S(\cdot)$ converges uniformly to $\mathcal{E}(\cdot)$. For a given g, γ, the dual objective function can be written as an expectation over their corresponding $f \in F$,

$$\mathcal{E}(g,\gamma) = \mathbb{E}_\alpha[f(X)],$$

and similarly for the empirical objective function

$$\mathcal{E}_S(g,\gamma) = \frac{1}{m}\sum_{t=1}^{m} f(X^t).$$

Then we can rewrite the supremum in (21) as:

$$\sup_{g,\gamma} |\mathcal{E}(g,\gamma) - \mathcal{E}_S(g,\gamma)| = \sup_{f \in F} \left| \mathbb{E}_\alpha[f(X)] - \frac{1}{m}\sum_{X \in S} f(X) \right| \tag{22}$$

Since $|f(X)| \leq (R\bar{x} + R)$ for all $f \in F, X \in \mathcal{X}$, it follows from Theorem 26.5 in [22] that with probability $1-\delta$,

$$\sup_{f \in F} \mathbb{E}_\alpha[f(X)] - \frac{1}{m}\sum_{X \in S} f(X) \leq 2\mathbb{E}_S\left[\text{Rad}_m(F \circ S)\right] + (R\bar{x} + R)\sqrt{\frac{2\log(2/\delta)}{m}} \tag{23}$$

and the same also holds by replacing F with $-F$. Here

$$\text{Rad}_m(F \circ S) := \mathbb{E}_\sigma\left[\frac{1}{m}\sup_f \sum_{j=1}^{m} \sigma_j f(X_j)\right]$$

is the standard definition of Rademacher complexity of the set $F \circ S$. Since σ_i are $i.i.d.$ Rademacher random variables, it is easy to see that $\text{Rad}_m(F \circ S) = \text{Rad}_m(-F \circ S)$. Therefore we can use a union bound to obtain that with probability $1-\delta$,

$$\sup_{f \in F} \left| \mathbb{E}_\alpha[f(X)] - \frac{1}{m}\sum_{X \in S} f(X) \right| \leq 2\mathbb{E}_S\left[\text{Rad}_m(F \circ S)\right] + (R\bar{x} + R)\sqrt{\frac{2\log(4/\delta)}{m}} \tag{24}$$

It remains to bound the Rademacher complexity of the $F \circ S$. To do so, we use tools from learning theory, and give the following bound on the fat-shattering dimension ([6]) of the hypothesis class F.

Lemma 1. *Under Assumption 1, F has ζ-fat-shattering dimension of at most $\frac{c_0(R\bar{x}+R)^2}{\zeta^2} n\log(n)$, where c_0 is some universal constant.*

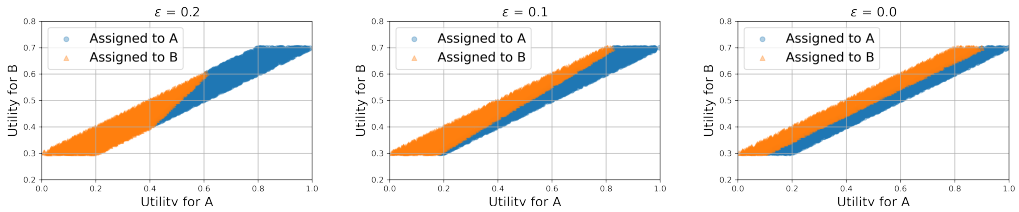

Figure 2: Allocation policy for artificial data under different envy constraints. From left to right: $\epsilon = 0.2, 0.1, 0.0$. When the envy constraint is loose (large ϵ), B envies A, since both agents prefer the items on the top right, but most of them are allocated to A. As the envy constraint tightens, the allocation boundary tilts in the direction that makes the allocations more even between the two agents.

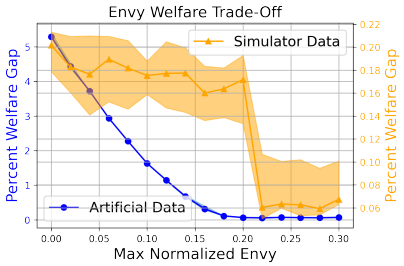

Figure 3: The trade-off curve between envy and welfare for both data-sets. The shaded region is between 25th and 75th percentile of the trials. The non-monotonicity in the plot for the simulator data is due to the stochasticity in the SGD algorithm.

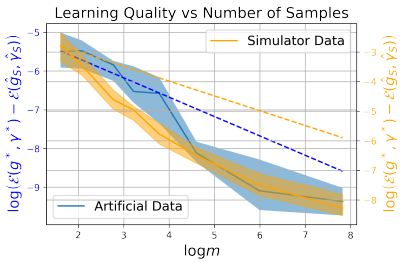

Figure 4: Approximation gap with respect to sample size. Both x and y axis are in log scale. The solid line is the median and the shaded region is between the 25th and 75th percentile. The dashed lines show what the theoretical $1/\sqrt{m}$ rate would look like.

Proof. Theorem 3 in [17] shows that $\text{fat}_\zeta(F_{min}) \leq \frac{c_0(R\bar{x}+R)^2}{\zeta^2} n \log n$. Since the shattering dimension is monotone in the size of the set, we are done.

□

The above bound on the fat-shattering dimension can be used to bound the covering number (see Definition 27.1 of [22]) of $F \circ S$. Theorem 1 from [19] states that

$$\mathcal{N}(\delta, F, ||\cdot||_2) \leq \left(\frac{2B}{\delta}\right)^{c_1 \text{fat}_{c_2\delta}(F)} \tag{25}$$

where B is a uniform bound on the absolute value of any $f \in F$. Let $B = (R\bar{x} + R)$, we have that

$$\begin{aligned}
&\text{Rad}_m(F \circ S) \\
&\leq \inf_{\delta'>0} \left\{ 4\delta' + 12 \int_{\delta'}^{B} \sqrt{\frac{\log \mathcal{N}(\delta, F, ||\cdot||_2)}{m}} d\delta \right\} \\
&\leq \inf_{\delta'>0} \left\{ 4\delta' + 12 \frac{\sqrt{c_1 c_0}}{c_2} B \sqrt{\frac{n \log n}{m}} \int_{\delta'}^{B} \sqrt{\log\left(\frac{2B}{\delta}\right)} d\delta \right\} \\
&= \inf_{\delta'>0} \left\{ 4\delta' + 8 \frac{\sqrt{c_1 c_0}}{c_2} B^2 \sqrt{\frac{n \log n}{m}} \left(\log(\frac{2B}{\delta'})^{\frac{3}{2}} + (\log 2)^{\frac{3}{2}} \right) \right\} \\
&= c' \sqrt{\frac{n \log n (\log m)^3}{m}}
\end{aligned}$$

Where in the first, second, and last step we used Dudley's chaining integral [24, 13], Lemma 1 and (25), and setting $\delta' = \frac{1}{\sqrt{m}}$ respectively. Plugging the above back to (24) and (21), we see that with

probability $1-\delta$,

$$\mathcal{E}(g^*, \gamma^*) - \mathcal{E}(\hat{g}_S, \hat{\gamma}_S) \leq c' \left(\sqrt{\frac{n \log n (\log m)^3}{m}} + \sqrt{\frac{1 \log \frac{1}{\delta}}{m}} \right).$$

Conversely, ignoring the log terms, m needs to be at most on the order of $\tilde{O}\left(\frac{n}{\epsilon^2}\right)$ in order for $\mathcal{E}(g^*, \gamma^*) - \mathcal{E}(\hat{g}_S, \hat{\gamma}_S)$ to be bounded by ϵ with high probability.

□

7 Experiments

We test our solution with both artificial data, and simulated data from a realistic simulator for blood donor matching developed by [18]. The artificial data contains two receivers, and their valuation distribution is a linearly transformed uniform distribution. This is to make visualization of the resulting allocation policy easier. The simulator data is based on geographical and population information from San Francisco, and contains 5 receivers. To set the envy budgets $\lambda \in \mathbb{R}^n$, we first decide on a constant $\epsilon \in \mathbb{R}_+$, and then multiply this by the target matching distribution $p^* \in \mathbb{R}^n$: $\lambda_{ij} = \epsilon p_i^* \, \forall i, j$. With this setup ϵ is a bound on the normalized envy for each receiver: $\frac{1}{p_i^*} Envy(i) \leq \epsilon \, \forall i$. Figure 2 illustrates how the allocation policy changes as we change ϵ. As the envy constraint tightens, the decision boundary tilts in the direction that split the "good" (items which both agents prefer) and "bad" (items which both agents dislike) items more evenly between the receivers.

Next we investigate the tradeoff between envy and social welfare by using SGD to compute approximately optimal allocations for varying ϵ. We plot the percent welfare gap (difference between the maximum welfare without envy constraints, and the welfare with envy constraints, divided by the former) with respect to realized, max normalized envy. Figure 3 shows the result. For the simulator data, the welfare gap is small even with a no-envy constraint, which means that aiming for envy free allocations might make sense. In the case of the artificial data however, paying 50% of the full price of fairness reduces 65% of the envy. In such settings, one might want to sacrifice some envy for better welfare.

These experiments also highlight when envy arises. When receivers' utilities are highly correlated, but one receiver has larger variance than others, that receiver receives almost all the good items (which results in large envy for other receivers), even though others value the items almost as much. In such cases, a small reduction in welfare can reduce a large amount of envy. This seems to be the case for the simulator data. On the other hand, if utilities are correlated, but only one receiver has very strong preferences, then allowing a small amount of envy can improve the welfare significantly.

Finally, in Figure 4 we investigate the quality of the empirical solutions as the sample size increases. It can been seen that the approximation gap decreases faster than the theoretical rate, confirming our sample complexity bound in Theorem 2.

8 Limitations and Future Directions

Although we believe that the model proposed here is natural, and captures the most salient aspects of some of the resource allocation problems in real life, any implementation of our proposed strategy in critical applications such as blood donation should be prefaced with more rigorous backtesting in order to minimize the risk of unintended consequences in application specific metrics not studied in this paper.

For future directions, one key property that we did not study in this paper is the problem of incentive compatibility. For our motivating application of blood donation, this is not an issue because online platforms such as Meta has proprietary models that can predict the matching quality between donor and receiver. This means that the platform observes the value of matchings without having to rely on the receivers to self-report. This is also true in many other online matching problems such as sponsored ads. However, in settings where the central planner relies on the receivers to self-report their valuations for each of the items, incentive compatibility becomes a crucial issue. We are excited about the potential of using Optimal Transport in fair-division, and plan on exploring the incentive issues in future work.

References

[1] Shipra Agrawal and Nikhil R Devanur. "Fast algorithms for online stochastic convex programming". In: *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. SIAM. 2014, pp. 1405–1424.

[2] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. "A dynamic near-optimal algorithm for online linear programming". In: *Operations Research* 62.4 (2014), pp. 876–890.

[3] Genevay Aude et al. "Stochastic optimization for large-scale optimal transport". In: *arXiv preprint arXiv:1605.08527* (2016).

[4] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. "Regularized online allocation problems: Fairness and beyond". In: *International Conference on Machine Learning*. PMLR. 2021, pp. 630–639.

[5] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. "The best of many worlds: Dual mirror descent for online allocation problems". In: *arXiv preprint arXiv:2011.10124* (2020).

[6] Peter L Bartlett, Philip M Long, and Robert C Williamson. "Fat-shattering and the learnability of real-valued functions". In: *journal of computer and system sciences* 52.3 (1996), pp. 434–452.

[7] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. "The price of fairness". In: *Operations research* 59.1 (2011), pp. 17–31.

[8] Ioannis Caragiannis et al. "The efficiency of fair division". In: *International Workshop on Internet and Network Economics*. Springer. 2009, pp. 475–482.

[9] Yiling Chen et al. "Truth, justice, and cake cutting". In: *Games and Economic Behavior* 77.1 (2013), pp. 284–297.

[10] Yuga J Cohler et al. "Optimal envy-free cake cutting". In: *Twenty-Fifth AAAI Conference on Artificial Intelligence*. 2011.

[11] Nikhil R Devanur and Thomas P Hayes. "The adwords problem: online keyword matching with budgeted bidders under random permutations". In: *Proceedings of the 10th ACM conference on Electronic commerce*. 2009, pp. 71–78.

[12] Nikhil R Devanur et al. "Near optimal online algorithms and fast approximation algorithms for resource allocation problems". In: *Journal of the ACM (JACM)* 66.1 (2019), pp. 1–41.

[13] Richard M Dudley. "The sizes of compact subsets of Hilbert space and continuity of Gaussian processes". In: *Journal of Functional Analysis* 1.3 (1967), pp. 290–330.

[14] Gagan Goel and Aranyak Mehta. "Online budgeted matching in random input models with applications to Adwords." In: *SODA*. Vol. 8. Citeseer. 2008, pp. 982–991.

[15] Nicholas Harvey. "Machine Learning Theory Lecture 17". In: (2018).

[16] Leonid Kantorovich. "On the transfer of masses (in russian)". In: *Doklady Akademii Nauk,* 37(2) (1942), pp. 227–229.

[17] Aryeh Kontorovich and Idan Attias. "Fat-shattering dimension of k-fold maxima". In: *arXiv preprint arXiv:2110.04763* (2021).

[18] Duncan C McElfresh et al. "Matching algorithms for blood donation". In: *Proceedings of the 21st ACM Conference on Economics and Computation*. 2020, pp. 463–464.

[19] Shahar Mendelson and Roman Vershynin. "Entropy and the combinatorial dimension". In: *Inventiones mathematicae* 152.1 (2003), pp. 37–55.

[20] Elchanan Mossel and Omer Tamuz. "Truthful fair division". In: *International Symposium on Algorithmic Game Theory*. Springer. 2010, pp. 288–299.

[21] Jack Robertson and William Webb. *Cake-cutting algorithms: Be fair if you can*. CRC Press, 1998.

[22] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

[23] Sean R Sinclair, Siddhartha Banerjee, and Christina Lee Yu. "Sequential Fair Allocation: Achieving the Optimal Envy-Efficiency Tradeoff Curve". In: *arXiv preprint arXiv:2105.05308* (2021).

[24] Karthik Sridharan. *Note on Refined Dudley Integral Covering Number Bound*. 2021. URL: https://www.cs.cornell.edu/~sridharan/dudley.pdf.

[25] Steven Yin et al. "Dominant Resource Fairness with Meta-Types". In: *arXiv preprint arXiv:2007.11961* (2020).

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 8
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 8
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Included as supplemental material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Appendix A.2
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Figure 3,4
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix A.2
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] Yes, this cited paper is this one: [18]
   (b) Did you mention the license of the assets? [Yes] See Appendix A.2
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Our code is provided as supplemental material.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] We used the simulator developed, and made public by the authors of a paper that we cite. There is no personal information in simulator data.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A] We did not crowdsource or conduct research with human subjects.
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]