

Fairness in Scarce Societal Resource Allocation: A Case Study in Homelessness Applications

Ashwin Kumar

Washington University in St. Louis
Saint Louis, MO, United States

William Yeoh

Washington University in St. Louis
Saint Louis, MO, United States

ABSTRACT

Societal resource allocation is a broadly studied field, where a limited quantity of resources is to be divided amongst a large population, like distributing limited beds among homeless people, or splitting food supplies between refugee camps. When group disparities exist, fairness becomes a critical issue because of the sensitive nature of the resources. Further, such systems may require repeated matching over time, thus making it an online problem. In this paper, we cast online scarce societal resource allocation as a multi-agent Markov Decision Process (MMDP) and demonstrate a simple incentives-based approach that can be used to improve fairness across a variety of trade-off weights. With the application domain of allocating homelessness services, we show our method’s efficacy on different group divisions and demonstrate competitive tradeoffs with overall system efficiency. Further, we show how only affecting a subset of the population can lead to better overall solutions. We provide a general discussion and results from our experiments to show how these methods compare against each other.

KEYWORDS

Fairness, Resource Allocation, Fair Division

1 INTRODUCTION

Many real-world situations demand the allocation of some restricted resources across a group of individuals in need. From allocation of food reserves to refugee shelters to distribution of homelessness resources to people needing them, to allocating interventions for child protection services, we are often faced with the need to allocate interventions and resources to individuals to maximize the benefit received to society. Allocating scarce societal resources is a challenging problem due to the heterogeneous needs of individuals receiving them. It is often difficult to satisfy every individual’s requirements to an equal extent, thus an efficient allocation aims to maximize some notion of aggregate utility. The utilitarian solution may also result in envy or a feeling of injustice for some individuals in the population, and this makes the study of fairness in these allocations important.

Prior work has looked at equitable allocation of resources across a given population using different notions of fairness, like serving the neediest people first, serving the most-helped first, or ensuring that the worst-off individual’s utility is maximized. In this work, we focus on the online setting, where resources and recipients arrive over time and batched decisions need to be made by a central authority. Instead of proposing a fair algorithm that optimizes some measure of fairness to generate an allocation, we propose an incentive based approach to allow for a continuous trade-off between fairness and efficiency. We begin by describing a general Multi-agent Markov Decision Process (MMDP) framework that

encapsulates general resource allocation problems. Then, using the group fairness notion of statistical parity, we show how we can move from an efficient allocation towards a fairer one through using the gradient information from the chosen fairness metric. We run experiments on a homelessness intervention allocation database, where the goal is to minimize total probability of re-entry into homelessness. Our experiments show a significant improvement in fairness is possible with minimal increase in overall probability of re-entry, especially when our incentives are applied only to a subset of the population.

2 RELATED WORK

Many real-world scenarios demand the division of resources (tasks) between multiple individuals (agents), each of whom derive different utilities for receiving them. From dividing a cake evenly [9] to efficiently dividing network bandwidth across multiple consumers [21], from allocating passengers to available taxis in ridesharing [26] to assigning homelessness interventions [19], the use-cases for resource allocation are diverse and far-reaching.

Efficient resource allocation has been a popular endeavour in the economics and computational optimization research community over the last several decades. The notion of *Social Good* frequently comes up in discussions about efficient resource allocation, where in addition to maximizing utility (the utilitarian approach), the decision-maker is also interested in enforcing fairness in how the resource is allocated to the agents [6, 8, 17]. There are a myriad of criteria for fairness, each defining fairness in different ways: minimizing variation among the population, maximizing the worst individual [5, 25] and allocating resources proportional to utility [18], being some examples. In some cases, fairness and utility can go hand in hand, while in (most) others, there exists a trade-off between fairness and utility. Especially with indivisible goods, a fair division may not exist, introducing the need for approximate fairness [22].

There exists a class of subproblems within resource allocation, where the matching (assignment of resources to agents) needs to be done in an online manner by a central decision-maker, with agents and/or resources able to re-enter the market over time. In this work, we call such problems *Temporal Resource Allocation Problems* (TRAPs). With uncertainty about future availability of resources, myopic decisions in the present are suboptimal in the long run (if they are even tractable to compute in the first place). In such systems, statistical approaches like machine learning and reinforcement learning are often used to evaluate “values” or expected utilities of present actions (assignments) to better guide decision making. We refer to such estimators as Value Function Approximators (VFAs) in this work.

Ensuring fairness in online matching is a much more complicated task, as the temporal dynamics engage in a tug-of-war with

local (temporal) fairness. Further, the issue of *Algorithmic Fairness* also creeps in with the use of VFAs, as their predictions may contain biases inherited from the data used to train them. There is a vast volume of work on algorithmic fairness, and it has been gaining recent popularity, with concepts like Statistical Parity (also called demographic parity) [12, 14], Equal Opportunity/Equalized Odds [15] and Predictive Value [10]. This area of research was famously spurred on after the COMPAS recidivism predictor was shown to be racially biased [3], bringing black-box algorithms used in decision-making under heavy scrutiny.

There has been some research into online group fairness in matching algorithms. Hosseini et al. [16] approach the task of creating an online approximate algorithm for class envy-freeness. However, they ignore the presence of utility estimators, and instead tackle fairness in a zero-one setting, where agents either want a resource or do not. Esmaili et al. [13] use both individual and group Rawlsian fairness criteria to improve two-sided fairness in online bipartite matching, and provide probabilistic methods to improve fairness. However, they consider a purely online setting with decisions being made in rounds, one decision each round. This differs significantly from our approach, where we batch the agents over a time window to have a locally optimal combinatorial solution, which introduces further complications to the optimization. Kumar et al. [20] deal with a much closer problem setting, with multi-agent batched resource allocation in the context of rideshare matching. We draw on the approach they presented to guide our system design.

3 PROBLEM FORMULATION

In this section, we describe the general structure of a TRAP, rooted in a constrained Multiagent MDP (MMDP) [7]. In an MDP, a decision-maker interacts with a stochastic environment, choosing an action a from a set of available actions A in the current state s . Performing the action leads to an observed reward r , as well as a transition to the next state s' , determined by a transition function $T(s, a, s')$.

A framework for TRAPs needs to capture the repeated interactions in a synchronous multi-agent setting, where the actions are allocations of indivisible resources to agents. We define the class of problems concretely below, drawing from Multiagent MDPs.

Definition 3.1 (TRAP). A temporal resource allocation problem is defined by the tuple $\langle \alpha, S, \{A_i\}_{i \in \alpha}, T, \{R_i\}_{i \in \alpha}, \gamma, c \rangle$, where:

- α is the set of all agents, each identified by $i \in \{1, 2, 3, \dots, n\}$;
- S is the set of global states;
- A_i is the set of actions for each agent i ;
- $T : S \times A_1 \times A_2 \times \dots \times A_n \times S \rightarrow [0, 1]$ is the joint transition function, encoding the probability of transition $T(s, a_1, a_2, \dots, a_n, s') = P(s' | s, a_1, a_2, \dots, a_n)$;
- $R_i : S \times A_i \rightarrow \mathbb{R}$ is the reward for agent i for a given action;
- γ is the discount factor that weighs the reward achieved in future time steps; and
- $c : A_1 \cup A_2 \cup \dots \cup A_n \rightarrow \mathbb{R}^K$ is a function that maps each action to its resource consumption.

We include separate rewards for each agent, akin to Multi-Objective MDPs (MOMDPs) [27]. In MOMDPs, the reward $R(s, a) \in$

\mathbb{R}^D is a vector, where each component is an objective to be optimized. This could be interpreted as two competing objectives (like fairness and efficiency), or as separate rewards for each agent. In a cooperative setting, the decision-maker's goal is to maximize total agent rewards over time.

A policy $\pi : S \rightarrow \mathcal{A}$ maps the global state to an allocation $\mathcal{A} = \{A_1 \times A_2 \times \dots \times A_n\}$, which represents the set of assigned actions for each agent. The joint action space even for a moderate number of agents is exponential in the size of α , so enumerating over all possible actions is intractable. Instead, agents actions are often assumed to be independent [11], allowing a simplification of the action space. In this system, each agent has its own associated state which it observes as a subset of the global state.

However, in resource allocation problems, agents cannot be truly separated. In this setting, an agent's action amounts to allocating to it a quantity of available resource. Since the overall availability of resources may be limited, this imposes a coupling between agents through the need to share resources. An agent's action is limited by the actions of other agents that require the same resource. This can be overcome by restricting the action space to only include joint actions that conform to the available resources. We formally include these constraints as part of the problem formulation for TRAPs. If K is the number of resources, each represented as $j \in \{1, 2, \dots, K\}$, the resource availability is written as $\mathcal{R} \in \mathbb{R}^K$. We thus say that each resource j has an availability \mathcal{R}_j , and this information is included in the global state. Further, each action a_i has an associated resource consumption $c(a_i) \in \mathbb{R}^K$.

For an allocation \mathcal{A} to be valid, the following constraints must be satisfied:

$$\begin{aligned} \sum_{a \in A_i, x_i(a) \in \{0,1\}} x_i(a) &= 1, \quad \forall i \in \alpha \quad (\text{Action Constraint}) \quad (1) \\ \sum_{a \in \mathcal{A}} c(a)_j &\leq \mathcal{R}_j, \quad \forall j \in \{1, \dots, K\} \quad (\text{Resource Constraint}) \quad (2) \end{aligned}$$

These constraints say each agent must be assigned exactly one action (Eq.1), and the combined resource consumption should not exceed the resource availability (Eq.2). Resources may be replenished after consumption, but their arrival rate is stochastic. Further, resource may be indivisible, in which case, additional integer constraints are added to each agent's actions. Agents' action sets may be padded with a null action to always ensure a feasible allocation.

In this work, agents represent the individuals looking for resources. In a scarce resource setting, number of agents is more than the available resources, $|\alpha| \geq K$.

3.1 Solving TRAPs Using Utility Approximation

The expected value of a policy can be written using the Bellman equation [4]:

$$V_\pi(s) = R(s, \pi(s)) + \gamma \sum_{s' \in S} T(s, \pi(s), s') V_\pi(s') \quad (3)$$

The optimal policy selects the action maximizing expected value. In this case, this means selecting a combination of agent actions that maximizes the cumulative action-value across all agents (for a cooperative system). Combinatorial optimization is often used to

find the best assignment subject to resource constraints, using techniques like linear programming or mixed-integer programming.

In many real-world scenarios with unknown transitions, infinite state spaces or large action spaces, computing the exact value function becomes impossible. Thus, practical approaches approximate this value function through a variety of statistical methods. Value Function Approximators (VFAs) try to estimate utility of taking an action, which guides the joint action selection at the current time step. Thus, we can write the action-value for an agent-allocation matching as follows:

$$Q(s_i, a_i) = R(s_i, a_i) + \gamma \sum_{s'_i \in S} T(s_i, a, s'_i) V(s'_i) \quad (4)$$

Given this, we can formulate the following optimization objective to maximize the cumulative action-values, where $x_i(a)$ is an indicator variable for action a being assigned to agent i :

$$\max_{x_i(a) \in \{0,1\}} \sum_{i \in \alpha} \sum_{a \in A_i} x_i(a) Q(i, a) \quad (5)$$

such that combining this objective with the action and resource constraints (Eqs.1-2) gives us a mixed-integer program that we can solve to find the optimal matching of agents to resources. This formulation is general and seen in many resource allocation problems [2, 19, 26].

3.2 Including Fairness in TRAPs

In this work, we tackle the problem of group fairness [29]. In a broad sense, group fairness aims to equalize some statistical measure over partitions of the population into groups. In our exploration, we look at groups defined both over the population of agents as well as over the population of resources. To include fairness, we can define a group metric $z : \{1, 2, \dots, g\} \rightarrow \mathbb{R}$ to map a group to its metric value. This value is updated based on the actions taken, after the transition to the successor state. Further, each agent or resource can be mapped to a group using a function $G : \alpha \cup \mathcal{R} \rightarrow \{1, 2, \dots, g\}$.

Each allocation \mathcal{A} has an effect of the distribution of group metric values $\mathbf{Z} = \{z_i\}_{i \in \{1,2,\dots,g\}}$, given by the function \mathcal{Z} . Let $\mathbf{Z}_{t+1} = \mathcal{Z}(\mathcal{A}_t^\pi, \mathbf{Z}_t)$ with policy π , and let $\mathcal{F}(\mathbf{Z})$ be a fairness metric over the group metrics. Our fairness objective is to minimize this fairness metric at a time horizon $t+k$, $\mathcal{F}(\mathbf{Z}_{t+k})$ by modifying the policy π . In a way, we are looking to minimize quality of service differences over groups over a time horizon, instead of improving individual-level allocation inequity at a single time step. The rationale behind this is that over a long time horizon, equalizing group-level treatment will also equalize the expected individual level treatment conditional on group membership.

4 SIMPLE INCENTIVES

In settings where the VFA is only accessible as a black box, one of the few options to enforce fairness is to post-process the value predictions to skew them in a fairness-aware manner. We draw inspiration from a post-processing method presented by Kumar et al. [20] called Simple Incentives. We describe the details of this method in the rest of this section.

We use the idea of *Statistical Parity* [1, 14] as our notion of fairness. Also called Demographic Parity, it requires that the expected value of a given metric z over a group g is the same as

$\bar{z} = \mathbb{E}_{g' \in G} [z(g')]$ the expected value of that same metric over all groups $g' \in G$:

$$z(g) = \bar{z} \quad (6)$$

$$|z - z(g)| \leq \epsilon \quad \forall g \in G \quad (7)$$

We can use a slack parameter ϵ to have a less strict requirement (Eq.7). While the goal of achieving parity is noble, achieving it through the matching of a single time step is rarely possible, especially when there is a large disparity. Instead, it is often better to look at amortized parity over a longer period of time [28].

To do this, we aim for a matching that “moves closer” towards parity with the goal of achieving parity in the near future. Towards that end, our framework uses variance $\text{var}(\mathbf{Z})$, where $\mathbf{Z} = \{z(g)\}_{g \in G}$ is the set of metric values for all groups, as a proxy measure for fairness and, at each time step, it takes a gradient step in the solution space, moving in the direction that minimizes variance.

If we assume that the average of the metric over all groups is stable (i.e., $\frac{\partial}{\partial \mathcal{A}} \bar{z} \simeq 0$, a reasonable assumption if a long enough history is included), then we can find an assignment for a modified score function that accounts for the gradient of the variance with respect to the assignment \mathcal{A} :

$$\begin{aligned} s'(s_i, a) &= Q(s_i, a) - \lambda \frac{\partial}{\partial \mathcal{A}} \text{var}(\mathbf{Z}) \\ &= Q(s_i, a) - \frac{1}{|\mathbf{Z}|} \lambda \frac{\partial}{\partial \mathcal{A}} \sum_{z_j \in \mathbf{Z}} (z_j - \bar{z})^2 \\ &= Q(s_i, a) + \frac{2}{|\mathbf{Z}|} \lambda \sum_{z_j \in \mathbf{Z}} (\bar{z} - z_j) \frac{\partial z_j}{\partial \mathcal{A}} \end{aligned} \quad (8)$$

where λ is a hyperparameter. The general form above for the second term is our *incentive* score: A constant (weight) multiplied by the disparity of group j , scaled by a derivative term. We show later that the sum can usually be simplified within the context of a given action a and the derivative can be approximated for specific metrics.

The “Simple Incentives” idea is that, for each group involved in an action, we provide them with an incentive (or penalty) proportional to how disadvantaged (or advantaged) their group has been historically. Given the recent abundance of black-box algorithms, we find this simplicity helpful from a transparency perspective, making it easy to explain to any stakeholder how this score is calculated.

5 HOMELESSNESS SERVICES

In this section, we instantiate a resource allocation problem on a homelessness prevention dataset [19]. Homelessness is a long-standing social issue that affects over half a million people annually in the US [23]. Congress allocates funds to support homeless people using different kinds of interventions, and individuals or families are assigned to these interventions on a case-by-case basis. Recent work [19, 24] looks at using a data-driven approach to improve the quality of these assignments, by estimating counterfactual probabilities of *re-entry* into homelessness. The objective is to create an alternative assignment that reduces the overall probability of re-entry across all people availing homeless services.

We formulate the homelessness-intervention allocation problem as a TRAP, represented by the tuple $\langle \alpha, S, \{A_i\}_{i \in \alpha}, T, R, \gamma, c \rangle$. In

this model, the interventions I are treated as the resources, with limited capacity C_i that may become available at future times. The agents $h \in \alpha$ correspond to the various households looking for interventions, where each household h is associated with a state S_h that captures some demographic information about the household, as well as their current status.

An action $a_h^i \in A_h$ corresponds to matching a household to an available intervention $i \in I$. $R(h, i)$ is the benefit of assigning households h to intervention $i \in I$. To avoid discriminating between individuals, we keep this a unit reward for all allocations. The utility of an assignment, as described in existing work [19] is just the probability of re-entry into homelessness when household h is matched to intervention i , $\Pr(h, i)$.

Thus, we can write the action-value for this problem as:

$$Q(a_h^i) = R(s_h, i) - \Pr(h, i) \quad (9)$$

$$= 1 - \Pr(h, i) \quad (10)$$

An allocation \mathcal{A} is a concatenation of interventions assigned to each household, with \mathcal{A}_h denoting the intervention allocated to household h . The consumption function $c(\mathcal{A}_h)$ returns a vector of size $|I|$, showing which intervention was used by an action. The optimization problem can then be written as follows, following Eq.5. $x_h(i)$ is an indicator variable denoting whether household h was matched to intervention i .

$$\max_{x_h(i) \in \{0,1\}} \sum_{h \in \alpha} \sum_{i \in I} x_h(i) Q(a_h^i) \quad \text{s.t.} \quad (11)$$

$$\sum_{i \in I, x_h(i) \in \{0,1\}} x_h(i) = 1, \quad \forall h \in \alpha \quad (12)$$

$$\sum_{a \in \mathcal{A}} c(a)_i \leq C_i, \quad \forall i \in I \quad (13)$$

Equations 12 and 13 show that each household must be allocated one intervention, and the total capacity for each intervention must not be violated. We assume that we have exactly as many interventions as households. Since the dataset (from [19]) contains the actual intervention that was assigned to each household, we can use that to construct the capacities of each intervention. The ILP we formulated here is equivalent to the one proposed by Kube et al. [19].

We also note that there is an online nature to this problem. The original dataset spans a few years, but in real life, we cannot wait for years before allocating resources to homeless people in need. Thus, we use information about the arrival dates of various households and batch them into groups over discrete time windows of length 30 days. Each time window, we accumulate the arrived households and solve the optimization problem stated above to allocate interventions. For our experiments here, we do not model the exit dates of the households, and thus, consumed interventions do not become available at a later date again.

5.1 Simple Incentives and Homelessness

The households in the previous setup may belong to various different groups, as described by features in their state description s_h . Let $G : \alpha \rightarrow \{1, 2, \dots, g\}$ be a general function that maps the group membership of household h to group j uniquely. Further, let z_j denote the historical average probability of re-entry for group j ,

Table 1: The effect of the fairness function $F(a_h^i)$ for $\beta > 0$.

| | Better-off group $\bar{z} - z(h) > 0$ | Worse-off group $\bar{z} - z(h) < 0$ |
|--|---|--|
| Bad action $\Pr(h, i) - z(h) > 0$ | Make this action seem less bad | Make this action seem worse |
| Good action $\Pr(h, i) - z(h) < 0$ | Make this action seem less good | Make this action seem better |

and \bar{z} denote the average of all group re-entry probabilities. Further, let $z(h) = z_{G(h)}$ map households to their group metric value. Then, it is desirable to have similar probabilities of re-entry across groups while having a low overall re-entry probability. To incorporate the incentive-based fairness function as described in Eq.8, we construct the following modified score function.

$$\begin{aligned} s'(a) &= Q(a) + \beta \sum_{z_j \in \mathbf{Z}} (\bar{z} - z_j) \frac{\partial z_j}{\partial \mathcal{A}} \\ &= Q(a) + \beta \sum_{z_j \in \mathbf{Z}} (\bar{z} - z_j) \frac{\partial z_j}{\partial a} \end{aligned} \quad (14)$$

Any action a serves to modify only the connected household's group probability. So, for all $z_j \neq z(h)$, the gradient term is zero. Thus, we can simplify Eq.14 as:

$$s'(a) = Q(a) + \beta (\bar{z} - z(h)) \frac{\partial z(h)}{\partial a} \quad (15)$$

The change in $z(h)$ depends on the history being considered. To not bias this by group size, we instead just approximate $\frac{\partial z(h)}{\partial a} = \Pr(h, i) - z(h)$. Any action that has a larger probability than the current group metric increases it, and vice-versa. The magnitude of change is also proportional to the difference, so this serves as a good substitute for the gradient.

We can write a fairness score $F(a)$ as follows:

$$\begin{aligned} F(a) &= \text{group_adv} \times \text{action_adv} \quad (16) \\ \text{group_adv} &= \bar{z} - z(h) \\ \text{action_adv} &= \Pr(h, i) - z(h) \end{aligned}$$

The sign of the group advantage tells us whether the group $G(h)$ is better than the average group or not, and the action advantage tells us whether the intervention i leads to a reduction (or increase) in the group metric.

Thus, the final form of the score is as follows:

$$s_\beta(a) = Q(a) + \beta F(a) \quad (17)$$

Here, β is a hyperparameter that controls the value of fairness. Table 1 shows the effects of this fairness score on various types of groups and actions. In general, this has the effect of making actions for worse-off groups seem much more consequential, thus incentivizing picking of a better option over a worse one, while it makes better-off groups' choices seem less consequential, so it is okay to trade-off more of their value to serve the worse off groups.

Additionally, we also explore the effect of only performing a subset of the score modifications.

- SI(+): We only use the bonus when it makes actions seem better than they are, i.e. when $F(a)$ is positive.

$$s_{\beta}(a) = Q(a) + \beta \max(F(a), 0) \quad (18)$$

- SI(-): We only use the bonus when it makes actions seem worse i.e. when $F(a)$ is negative.

$$s_{\beta}(a) = Q(a) + \beta \min(F(a), 0) \quad (19)$$

We expect both these approaches to have different effects, as they make different combinations of actions more appealing to the decision maker (ILP). As a reminder, the ILP maximizes the action-values. Since the action value contains the negative of the probability of re-entry, it minimizes the cumulative probability of re-entry across all assignments.

6 EXPERIMENTS

We used the counterfactual probabilities generated using Bayesian Additive Regression Trees (BART) as presented by Kube et al. [19] as our estimates of re-entry probability for each household when matched to one of four interventions: Emergency Shelter (ES), Transitional Housing (TH), Rapid Re-housing (RRH) and Homelessness Prevention measures (Prev). Descriptions of these interventions can be found in the original paper [19]. Each household has a heterogeneous order of utility over the available interventions. In total, there were 13,940 households. There were differing counts of each intervention available: Prev: 6202, ES: 4441, TH: 2451, RRH: 846. As stated earlier, we assumed each intervention slot could only be assigned once and would not re-enter the system, and similarly, each household was also considered for entry only once.

Each household was associated with close to 50 features. Of these, 38 features had at least two and at most twenty unique groups, excluding groups with less than 50 members. We considered each of these groups as potential groups of interest, and ran experiments to see if our method could improve fairness on arbitrary group divisions.

We ran experiments for a variety of β values in a logarithmic grid search, starting at $\beta = 10$ and ending at $\beta = 10000$. For each feature used as a group identifier, we ran 3 experiments, one with the basic SI, one with SI(+) and one with SI(-).

As a metric for fairness, we use the Gini coefficient, a popular metric used to measure income inequality in a population. We evaluate $Gini(\mathbf{Z})$ over all group metrics. A Gini coefficient of 0 implies perfect equality, so we want to minimize this metric. Our metric for efficiency is the overall probability of re-entry, which is just the aggregate re-entry probability $\Pr(h, i)$ for all allocated household-intervention pairs. We want this to be low as well.

In our experiments, we compare our fair solutions with the fairness incentives to a baseline ILP matching (Eq.11). To do this, we report the Price of Fairness (PoF) and the Benefit of Fairness (BoF). The PoF compares the loss in efficiency vis-a-vis the re-entry probability. Let PRE_{opt} be the re-entry probability with the optimal ILP solution, and PRE_{SI} be the re-entry probability with one of our fairness methods. Then,

$$PoF = \frac{PRE_{SI}}{PRE_{opt}} \quad (20)$$

Distribution of HousingStatusAtEntry




Figure 1: The distribution of average probabilities across different groups by solution method for the feature "HousingStatusAtEntry".

Here, we expect suboptimal solutions to have $PoF > 1$. Similarly, if $Gini_{opt}$ results from the optimal ILP and $Gini_{SI}$ results from our methods,

$$BoF = \frac{Gini_{SI}}{Gini_{opt}} \quad (21)$$

Here, we expect our solutions to have $BoF < 1$.

6.1 Evaluation Results

For each feature used as group information, we report the BoF for the best $Gini(\mathbf{Z})$ obtained by our method for any β value, while maintaining the PoF to be below 1.05. We also report the corresponding β value that led to that assignment. The results are presented in Table 2.

We observe that we are generally able to significantly improve the fairness using any of our methods for most group divisions. There are some instances (like PrimaryRace) where using SI-based fairness doesn't improve the distribution at all, but barring these exceptions, we observe a few trends:

- SI(-) performs really well for most features. Nearly two-thirds of the time (25/38), it gets the best BoF with low PoF.
- SI is sometimes able to do really well (9/38)
- SI(+) is the poorest performing method.

This suggests that only improving the probabilities of actions is not a good strategy, while typically, only reducing the probabilities (making bad options worse for worse-off groups and good options seem less good for advantaged groups) is successful in allowing for better fairness over repeated allocations over time. We also observe that there is a high variability in this value across features, indicating that different features have different responses to the trade-off weight.

6.2 Empirical Observations: Selected Results

Here, we pick one representative group feature and analyze it in detail, to provide readers with a better understanding of our results. We select the feature "HousingStatusAtEntry", which captures the current housing status of a household before intervention. Using this feature to inform group division, we get the following groups: 1: Homeless (1156 households), 2: at imminent risk of losing housing (3551 households), 3: at risk of homelessness (1242 households), 4:

Table 2: Fairness improvement for group divisions based on various features. For each method, we show the best Benefit of Fairness (BoF) achieved while restricting the Price of Fairness (PoF) to be below 1.05.

| Group | Original | | SI | | | SI(+) | | | SI(-) | | |
|--|----------|--------|------|--------|---------------|-------|--------|---------------|-------|--------|---------------|
| | Prob | Gini | Beta | PoF | BoF | Beta | PoF | BoF | Beta | PoF | BoF |
| PrimaryRace | 0.2443 | 0.1114 | 0 | - | - | 0 | - | - | 25 | 1.0021 | 0.9982 |
| Gender | 0.2443 | 0.0184 | 100 | 1.0129 | 0.8723 | 500 | 1.0457 | 0.3977 | 2500 | 1.0227 | 0.0421 |
| Ethnicity | 0.2443 | 0.1752 | 25 | 1.0025 | 0.7231 | 25 | 1.0023 | 0.7210 | 1000 | 1.0031 | 0.8536 |
| VeteranStatus | 0.2443 | 0.0579 | 100 | 1.0066 | 0.9405 | 100 | 1.0431 | 0.6369 | 750 | 1.0146 | 0.6500 |
| DisablingCondition | 0.2443 | 0.1223 | 25 | 1.0489 | 0.9288 | 25 | 1.0231 | 0.9400 | 75 | 1.0484 | 0.8764 |
| HousingStatusAtEntry | 0.2443 | 0.1939 | 75 | 1.0410 | 0.5712 | 100 | 1.0436 | 0.6977 | 2500 | 1.0390 | 0.5273 |
| HUDChronicHomeless | 0.2443 | 0.0941 | 25 | 1.0288 | 0.8808 | 25 | 1.0199 | 0.8884 | 100 | 1.0330 | 0.5964 |
| PhysicalDisability | 0.2443 | 0.0580 | 250 | 1.0497 | 0.4089 | 250 | 1.0115 | 0.7993 | 10000 | 1.0475 | 0.5091 |
| ReceivePhysicalDisabilityServices | 0.2443 | 0.0372 | 100 | 1.0087 | 0.9510 | 250 | 1.0381 | 0.8608 | 2500 | 1.0204 | 0.6670 |
| HasDevelopmentalDisability | 0.2443 | 0.1386 | 10 | 1.0000 | 0.9988 | 10 | 1.0004 | 0.9983 | 1000 | 1.0476 | 0.4891 |
| ReceiveDevelopmentalDisabilityServices | 0.2443 | 0.0963 | 0 | - | - | 25 | 1.0437 | 0.7385 | 7500 | 1.0449 | 0.3466 |
| HasChronicHealthCondition | 0.2443 | 0.0264 | 0 | - | - | 250 | 1.0445 | 0.8548 | 2500 | 1.0373 | 0.6426 |
| ReceiveChronicHealthServices | 0.2443 | 0.0051 | 0 | - | - | 0 | - | - | 0 | - | - |
| HasHIV/AIDS | 0.2443 | 0.0372 | 1000 | 1.0132 | 0.4644 | 100 | 1.0103 | 0.6428 | 1000 | 0.9990 | 0.6677 |
| HasMentalHealthProblem | 0.2443 | 0.0863 | 50 | 1.0361 | 0.8037 | 75 | 1.0474 | 0.7868 | 500 | 1.0448 | 0.5828 |
| ReceiveMentalHealthServices | 0.2443 | 0.0793 | 50 | 1.0345 | 0.8205 | 50 | 1.0209 | 0.873 | 250 | 1.0379 | 0.6614 |
| HasSubstanceAbuseProblem | 0.2443 | 0.0940 | 50 | 1.0468 | 0.6092 | 75 | 1.0472 | 0.822 | 10000 | 1.0087 | 0.8402 |
| ReceiveSubstanceAbuseServices | 0.2443 | 0.0196 | 2500 | 1.0371 | 0.0899 | 7500 | 1.0422 | 0.4135 | 10000 | 1.0118 | 0.3769 |
| DomesticViolenceSurvivor | 0.2443 | 0.0659 | 50 | 1.0458 | 0.5524 | 25 | 1 | 0.9999 | 25 | - | - |
| proj_type_ent | 0.2443 | 0.1210 | 50 | 1.0488 | 0.6118 | 25 | 1.0285 | 0.8845 | 750 | 1.0471 | 0.5172 |
| proj_type_exit | 0.2443 | 0.1026 | 25 | 1.0427 | 0.6990 | 25 | 1.0354 | 0.8607 | 250 | 1.0446 | 0.5511 |
| numProj_before_exit | 0.2443 | 0.0449 | 0 | - | - | 0 | - | - | 250 | 1.0432 | 0.6672 |
| reentered | 0.2443 | 0.0981 | 50 | 1.0436 | 0.5927 | 75 | 1.0445 | 0.6915 | 1000 | 1.0332 | 0.4766 |
| reenteredNotStable | 0.2443 | 0.0739 | 75 | 1.0466 | 0.3994 | 100 | 1.0438 | 0.6735 | 1000 | 1.0285 | 0.4562 |
| reenteredFedDef | 0.2443 | 0.0982 | 50 | 1.0436 | 0.5930 | 75 | 1.0446 | 0.6914 | 1000 | 1.0332 | 0.4770 |
| reenterType | 0.2443 | 0.0891 | 25 | 1.0250 | 0.8621 | 50 | 1.0417 | 0.7413 | 1000 | 1.0489 | 0.4649 |
| HousingStatusNextCall | 0.2443 | 0.1419 | 25 | 1.0245 | 0.6551 | 50 | 1.0489 | 0.7788 | 750 | 1.0246 | 0.6690 |
| SpousePresent | 0.2443 | 0.0947 | 500 | 1.0415 | 0.4416 | 1000 | 1.0367 | 0.4943 | 5000 | 1.0436 | 0.4001 |
| Children | 0.2443 | 0.1871 | 100 | 1.0466 | 0.7330 | 100 | 1.0277 | 0.7512 | 2500 | 1.0339 | 0.9339 |
| Children0_2 | 0.2443 | 0.1977 | 50 | 1.0444 | 0.9373 | 100 | 1.0492 | 0.9192 | 250 | 1.0415 | 0.9162 |
| Children3_5 | 0.2443 | 0.0924 | 0 | - | - | 0 | - | - | 0 | - | - |
| Children6_10 | 0.2443 | 0.1434 | 25 | 1.0002 | 0.9638 | 25 | 1.0003 | 0.9648 | 500 | 1.0122 | 0.9558 |
| Children11_14 | 0.2443 | 0.1550 | 75 | 1.0024 | 0.9159 | 100 | 1.0269 | 0.9013 | 1000 | 1.0421 | 0.8860 |
| Children15_17 | 0.2443 | 0.2197 | 75 | 1.0032 | 0.8554 | 750 | 1.0136 | 0.8590 | 10000 | 1.003 | 0.8450 |
| num_members | 0.2443 | 0.1543 | 500 | 1.0487 | 0.7437 | 250 | 1.0440 | 0.7509 | 500 | 1.0109 | 0.8111 |
| UnrelatedChildren | 0.2443 | 0.1288 | 25 | 1.0138 | 0.7532 | 50 | 1.0341 | 0.7388 | 250 | 1.0422 | 0.6620 |
| UnrelatedAdults | 0.2443 | 0.2619 | 500 | 1.0175 | 0.8881 | 100 | 1.0053 | 0.9073 | 10000 | 1.0246 | 0.8912 |
| proj | 0.2443 | 0.1209 | 50 | 1.0487 | 0.6081 | 25 | 1.0285 | 0.8838 | 750 | 1.0473 | 0.5140 |

stably housed (402 households), 8: unknown status (6326 households) and -1: missing information (1263 households).

Figure 1 shows the distribution of average groupwise probabilities of re-entry for the four treatments (ILP, SI, SI(+), SI(-)) of the corresponding row in Table 2. We can see that for all SI methods, the group probabilities are shifted towards the mean. Empirically, this means we see an improvement in fairness, while the overall probability of re-entry is bounded to be within 5% of the original. This visually shows the intended effect of our incentive score: to move group scores closer to the mean.

Looking at Figure 2, we can more directly compare the performance of the three SI variants when β is varied. With increasing β , we see an improvement in fairness, but at a cost to overall utility.

Qualitatively, we see that the SI(-) method pareto dominates the others. For larger BoF ratios, we see that it is possible to improve fairness without significant increase in PoF. However, if much lower Gini values are required, SI seems to be the better approach, achieving much lower BoF albeit at a higher cost. On the other hand, SI(+) still seems to improve fairness, but at a much higher PoF. The general trend seems to favor SI and SI(-) over SI(+), as we observe similar behavior for most other features.

We note that the trends observed are not universal, and there are various features that, when used as group membership, do not work well with our framework to improve fairness. We hypothesize that the relationship between features and the utility estimation process plays a crucial role in deciding which method works well




Figure 2: The trade-offs in PoF and BoF for the feature “HousingStatusAtEntry”. The optimal point is to the bottom-left.

for that feature. Further experiments are needed to tease out the exact cause of the variable behavior.

7 CONCLUSION

In this work, we presented a general framework for modelling resource allocation problems with repeated matching, and using demographic parity as a starting point, derived a simple incentive-based scheme for continually improving fairness in an online setting. We showed three variants of our framework, and showed their efficacy through evaluation on a real-world dataset. We found interesting patterns in the behavior of each of the methods, which leave questions open for exploration in future work.

8 ETHICS STATEMENT

In this work, we take the data provided in our source material as is, to construct a matching problem and demonstrate a way of improving fairness if the value estimates are correctly identified. We do not claim that this approach will work in practice, as the feasibility of the matches may have real-world considerations not captured within the dataset. We only aim to provide a discussion for a possible method of improving group fairness, that can presents decision makers with some insight on similar online resource allocation problems.

REFERENCES

- [1] Alekh Agarwal, Miroslav Dudík, and Zhiwei Steven Wu. 2019. Fair Regression: Quantitative Definitions and Reduction-based Algorithms. In *International Conference on Machine Learning*. 120–129.
- [2] Javier Alonso-Mora, Samitha Samaranyake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. 2017. On-demand High-capacity Ride-sharing via Dynamic Trip-vehicle Assignment. *Proceedings of the National Academy of Sciences* 114 (2017), 462–467.
- [3] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine Bias. In *Ethics of Data and Analytics*. 254–264.
- [4] Richard Bellman. 1957. A Markovian Decision Process. *Journal of Mathematics and Mechanics* 6 (1957), 679–684.
- [5] Dimitri Bertsekas and Robert Gallager. 2021. *Data Networks*. Athena Scientific.
- [6] Hamoud S Bin-Obaid and Theodore B Trafalis. 2020. Fairness in Resource Allocation: Foundation and Applications. In *International Conference on Network Analysis*. 3–18.
- [7] Craig Boutilier. 1996. Planning, Learning and Coordination in Multiagent Decision Processes. In *Conference on Theoretical Aspects of Rationality and Knowledge*. 195–210.
- [8] Sylvain Bouveret, Yann Chevaleyre, Nicolas Maudet, and Hervé Moulin. 2016. *Fair Allocation of Indivisible Goods*. Cambridge University Press, 284–310.
- [9] Steven J Brams and Alan D Taylor. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- [10] Alexandra Chouldechova. 2017. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data* 5 (2017), 153–163.
- [11] Frits De Nijs, Erwin Walraven, Mathijs De Weerd, and Matthijs Spaan. 2021. Constrained Multiagent Markov Decision Processes: A Taxonomy of Problems and Algorithms. *Journal of Artificial Intelligence Research* 70 (2021), 955–1001.
- [12] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness Through Awareness. In *Innovations in Theoretical Computer Science Conference*. 214–226.
- [13] Seyed A Esmaili, Sharmila Dupplala, Vedant Nanda, Aravind Srinivasan, and John P Dickerson. 2022. Rawlsian Fairness in Online Bipartite Matching: Two-sided, Group, and Individual. *arXiv preprint arXiv:2201.06021* (2022).
- [14] Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and Removing Disparate Impact. In *ACM International Conference on Knowledge Discovery & Data Mining*. 259–268.
- [15] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*.
- [16] Hadi Hosseini, Zhiyi Huang, Ayumi Igarashi, and Nisarg Shah. 2022. Class Fairness in Online Matching. *arXiv preprint arXiv:2203.03751* (2022).
- [17] Shahin Jabbari. 2019. *Fairness and Feedback in Learning and Games*. Ph.D. Dissertation. University of Pennsylvania.
- [18] Frank P Kelly, Aman K Maulloo, and David Kim Hong Tan. 1998. Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability. *Journal of the Operational Research society* 49 (1998), 237–252.
- [19] Amanda Kube, Sanmay Das, and Patrick J Fowler. 2019. Allocating Interventions Based on Predicted Outcomes: A Case Study on Homelessness Services. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 622–629.
- [20] Ashwin Kumar, Yevgeniy Vorobeychik, and William Yeoh. 2023. Using Simple Incentives to Improve Two-Sided Fairness in Ridesharing Systems. In *International Conference on Automated Planning and Scheduling*.
- [21] Tian Lan, David Kao, Mung Chiang, and Ashutosh Sabharwal. 2010. *An Axiomatic Theory of Fairness in Network Resource Allocation*. IEEE.
- [22] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *ACM Conference on Electronic Commerce*. 125–131.
- [23] National Alliance to End Homelessness. 2022. State of Homelessness Archive. <https://endhomelessness.org/resource/archived-state-of-homelessness/>. Accessed: 2023-02-10.
- [24] Quan Nguyen, Sanmay Das, and Roman Garnett. 2021. Scarce Societal Resource Allocation and the Price of (Local) Justice. In *AAAI Conference on Artificial Intelligence*. 5628–5636.
- [25] Bozidar Radunovic and Jean-Yves Le Boudec. 2007. A Unified Framework for Max-Min and Min-Max Fairness with Applications. *IEEE/ACM Transactions on Networking* 15 (2007), 1073–1083.
- [26] Sanket Shah, Meghna Lowalekar, and Pradeep Varakantham. 2020. Neural Approximate Dynamic Programming for On-Demand Ride-Pooling. In *AAAI Conference on Artificial Intelligence*. 507–515.
- [27] Umer Siddique, Paul Weng, and Matthieu Zimmer. 2020. Learning Fair Policies in Multi-objective (Deep) Reinforcement Learning with Average and Discounted Rewards. In *International Conference on Machine Learning*. 8905–8915.
- [28] Tom Sühr, Asia J. Biega, Meike Zehlike, Krishna P. Gummadi, and Abhijnan Chakraborty. 2019. Two-Sided Fairness for Repeated Matchings in Two-Sided Markets. In *ACM International Conference on Knowledge Discovery & Data Mining*. 3082–3092.
- [29] Edwin Wenink. 2021. AI Ethics Group Fairness Taxonomy. <https://edwinwenink.github.io/ai-ethics-tool-landscape/fairness/group-fairness/>. Accessed: 2023-02-15.