Trade-offs between Group Fairness Metrics in Societal Resource Allocation

Tasfia Mashiat
George Mason University
Fairfax, VA, USA
tmashiat@gmu.edu

Xavier Gitiaux
George Mason University
Fairfax, VA, USA
xgitiaux@gmu.edu

Huzefa Rangwala
George Mason University
Fairfax, VA, USA
rangwala@gmu.edu

Patrick Fowler
Washington University in Saint Louis
St. Louis, MO, USA
pjfowler@wustl.edu

Sanmay Das
George Mason University
Fairfax, VA, USA
sanmay@gmu.edu

ABSTRACT

We consider social resource allocations that deliver an array of scarce supports to a diverse population. Such allocations pervade social service delivery, such as provision of homeless services and assignment of refugees to cities, among others. At issue is whether allocations are fair across sociodemographic groups and intersectional identities. Our paper shows that necessary trade-offs exist for fairness in the context of scarcity; many reasonable definitions of equitable outcomes cannot hold simultaneously except under stringent conditions. For example, defining fairness in terms of improvement over a baseline inherently conflicts with defining fairness in terms of loss compared with the best possible outcome. Moreover, we demonstrate that the fairness trade-offs stem from heterogeneity across groups in intervention responses. Administrative records on homeless service delivery offer a real-world example. Building on prior work, we measure utilities for each household as the probability of reentry into homeless services if given three homeless services. Heterogeneity in utility distributions (conditional on received services) for several sociodemographic groups (e.g. single women with children versus without children) generates divergence across fairness metrics. We argue that such heterogeneity, and thus, fairness trade-offs, pervade many social policy contexts.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; • **General and reference** → **Metrics**; **Evaluation**.

KEYWORDS

Resource allocation, algorithmic fairness, fairness metrics

ACM Reference Format:
Tasfia Mashiat, Xavier Gitiaux, Huzefa Rangwala, Patrick Fowler, and Sanmay Das. 2022. Trade-offs between Group Fairness Metrics in Societal Resource Allocation. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), June 21–24, 2022, Seoul, Republic of Korea.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3531146.3533171

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

ACM ISBN 978-1-4503-9352-2/22/06.
https://doi.org/10.1145/3531146.3533171

1 INTRODUCTION

Many social interventions that allocate resources to individuals are challenging because individuals have heterogeneous utilities. Thus, the design and analysis of allocation policies for social interventions in terms of efficiency and fairness is critical [44], as seen in many domains including child protection (e.g. [12]), healthcare (e.g. [47]), and homeless services (e.g [9, 36]). A particular concern for the use of machine learning posits that the tools systematically disfavor some sociodemographic or intersectional groups (see [13] for a review). For example, a growing body of work has documented racial disparities in credit lending, recidivism risk assessment [42], education [27], healthcare [41], and policing [21]. In this paper, we explore how to measure these potential disparities in the context of allocating resources given a limited budget. The literature on fair resource allocation has typically come from the areas of fair division and cooperative game theory. In that literature, one typically thinks of individuals as having preferences, and tries to define measures of fairness and allocation mechanisms that demonstrate these properties with respect to individual preferences. Recent notions of group fairness coming from the fair division line of literature strengthen the requirements for individual fairness [15] and are thus too strong for situations of scarce resource allocation, where allocations by definition must be unfavorable to some individuals.

So how should one measure fairness across groups in the allocation of scarce societal resources, where decisions often are made on the basis of multiple criteria? To ground our considerations in a specific case, consider homelessness service provision, where federal policy makes serving the most vulnerable an explicit goal, and at the same time, the effectiveness of services is measured by returns to homelessness among those served [1]. Such examples motivate us to consider how different notions of what role social services should play lead to different conclusions about the fairness of potential allocations across demographic groups.

For example, we could analyze how much better off members of a group are compared with how well they would have done under some minimal baseline allocation, or we could look at how much worse-off members of a group are than they would have been under the allocations that serve them the best. Fairness could then be

defined as equitable performance of groups according to these measures, and indeed, the existing literature on fair allocation of both divisible and indivisible resources has looked at measures along both of directions, of improvement (or gain) and regret (or shortfall). Although both are reasonable definitions of a fair allocation, we consider two important factors that arise in many real-world problems. First, instead of the problem simply focusing on a set of identical resources that need to be allocated amongst agents, there is often a whole set of different interventions, each with capacity constraints (for example, different types of homelessness resources or different cities that refugees can be matched to). Second, individuals may respond heterogeneously to the different interventions (for example, homeless individuals with disabilities may benefit disproportionately from intensive housing supports, or refugees may assimilate and find jobs more easily in places where there is already a substantial population from their place of origin).

We show that when there is a multiplicity of possible services, and groups are heterogeneous in the distributions of utilities they receive from different services, it becomes impossible to satisfy simultaneously improvement and regret oriented definitions of group fairness. Even more dramatically, an allocation policy that appears to favor one group according to improvement fairness can favor another group according to regret fairness. The results yield insights on inherent trade-offs that policymakers face when attempting to achieve a fairness objective. How we measure improvement or regret also matters when assessing the fairness of an allocation policy. For example, we could measure improvement by the ratio of realized utility over baseline utility (a multiplicative measure); or by the difference between realized utility and baseline utility (an additive measure). Depending on the application, it is not always clear which of these additive or multiplicative normalizations makes more sense. We establish, in a stylized framework, that fairness in terms of additive normalization and fairness in terms of multiplicative normalization cannot hold simultaneously except when the distribution of individual responses to different allocations is similar across demographic groups.

These trade-offs are not theoretical corner-cases and have substantive implications for social policy. We use administrative data from a regional homeless system to explore the fairness of a capacitated assignment of community-based services that address housing needs. Services include transitional housing, rapid rehousing, and emergency shelter; three programs that vary in intensity and availability. We measure the utility of a service to a household as the probability estimated in prior work by [36] that the household would make a successful exit from homelessness given the delivery of that service. We first document significant differences in utility distributions across different groups (e.g., disabled versus not disabled households, families with children versus households without children, single females with versus without children). We then confirm our theoretical results that the differences in utility distributions across groups generate trade-offs when assessing the fairness of an allocation. For example, we consider the original allocation as recorded in the administrative data and we find that improvement and regret disagree on whether the policy favors households with or without children, as well as other groups.

In addition to contributing to our understanding of how the definition and measurement of fairness is affected by heterogeneity in how members of different groups may respond to interventions, these findings can inform practice in homeless and social services that allocate scarce resources across diverse populations. Policies frequently attempt to maximize public welfare by targeting available supports towards heterogeneous groups based on competing notions of fairness (e.g., vulnerability, efficiency, equality). Understanding the fairness trade-offs and measurement sensitivity allows for more intentional policy-making and better evaluation.

2 RELATED WORK

2.1 Group Fairness

Prior research has led to many definitions of fairness to compare algorithmic outcomes across demographic groups. Popular definitions include statistical parity [19], equalized odds and opportunity [28]. However, these definitions only apply to binary settings and implicitly assume that the utility of an individual is equal to one when the algorithm's outcome is one and equal to zero otherwise. Few papers consider more general definitions of utilities [29]. In this paper, we argue as in [32] that in many societal applications of machine learning, utilities are heterogeneous across individuals and that this heterogeneity could be systematic across demographic groups.

The fair division literature offers a framework to compare utilities across individuals. Envy-freeness, proportionality or equitability [10] are common utility-based definitions of a fair allocation of goods. The literature strengthens these notions of fairness to control for envy-freeness to arbitrary segments of the population [6, 15]. In this paper, we focus on notions of group equitability that vary by their normalization, but leaves it for future research to explore the role of normalization on group envy-freeness.

A standard assumption in the fair division literature is that utilities, although heterogeneous, are unit-normalized [5]. The rationale for unit-normalization is that it allows one to make more reasonable interpersonal comparisons of utility by converting all utilities to a common scale. Unit-normalization implies that the maximum utility gain is equal to one for all individuals [5]. Our notions of shortfall or regret rely on a similar assumption, which is reasonable in many settings (e.g. voting [8]). However, we argue that other reasonable choices of normalization are possible and more relevant in different types of allocation problems. For example, in the case of homeless services delivery, a policymaker would want to account for the fact that families with children have on average more to gain from rapid rehousing programs [43]. In this case, our measures of improvement and gain, which normalize by comparison with the worst utility that an individual can expect from an allocation, are also reasonable notions of fairness. This paper relates closely to the work of [32], who introduce utility-based notions of group fairness for classification problems. However, they assume away the need to normalize utilities to a similar scale / support. One of our contributions is to show that different normalization approaches can lead to conflicting assessments of the fairness of an allocation policy.

2.2 Impossibility Results

The binary outcome setting admits some fundamental impossibility results [11, 35]. Except under very restrictive conditions, it is

impossible for a classifier to simultaneously equalize false positive rates and false negative rates across groups and also guarantee that predictions are calibrated within each group. Kleinberg et al. [35] show that the impossibility emerges whenever demographic groups differ systematically in the distribution of features used by the classifier as inputs. In this paper, we demonstrate new impossibility results in the case of utility-based notions of fairness. As in [35], we obtain a paradox where fairness guarantees that seem to share the same objective – that the allocation of resources will be as effective for all demographic groups – are nonetheless incompatible.

Our results on the incompatibility of different fairness principles is also reminiscent of Arrow's impossibility theorem [3]. In the presence of heterogeneous preferences, there is no way to aggregate individual preferences into a social welfare function that would satisfy unanimity, non-dictatorship and informational parsimony. The theory of fair allocation [25, 46] that selects a subset of policies on basis of their fairness and efficiency obtains possibility results by relaxing informational parsimony [24]. However, in this paper, we show that we cannot avoid negative results when notions of fairness based on different normalizations have to hold simultaneously.

2.3 Algorithmic Allocation of Societal Resources

There has been recent interest in the specific setting where scarce resources that are collective or societal are algorithmically allocated by a centralized institution to individual members of society (see [17] for a recent review). The design of algorithmic approaches has typically focused on increasing the efficiency of social interventions, including kidney exchange [37, 45], housing assistance [36, 38], HIV awareness campaigns [47] and refugee resettlement [18]. In this paper, we investigate how to assess the fairness of resulting allocations. Empirically, we find evidence of our impossibility results in the context of capacitated one-sided matching, which involve a set of services with fixed capacities, a set of agents with heterogeneous preference orderings (see e.g. [38] for an application to the house allocation problem) and a social worker that assigns a service to each agent.

3 INHERENT FAIRNESS TRADE-OFFS IN RESOURCE ALLOCATION

In this section we describe our theoretical framework, first defining the problems we are concerned with, and then outlining both general and illustrative results on inherent group fairness trade-offs in the allocation of scarce resources.

3.1 Setting

We consider K services, with maximum capacities c_k for $k \in \{1, ..., K\}$, and N individuals $i = 1, ..., N$.[1] We can thus describe individuals by their utility vector $\mathbf{u} = (u_1, ..., u_K)$ over each program k and their sensitive attribute $s \in \mathcal{S}$. $\mathcal{S}$ describes the set of groups for which we want to study the fairness of service allocation. For ease of exposition, we assume that group characteristics are binary and $\mathcal{S} = \{0, 1\}$; however, our results readily extend to more complex definitions of groups, and the empirical section will show that our results hold for intersectional groups. We denote by N_s the number of individuals with sensitive attribute $s = 0, 1$.

For each individual $\mathbf{u}$, we denote by $u^{\min}$ the utility derived from receiving the least beneficial program: $u^{\min} = \min\{u_k | k = 1, .., K\}$. We denote by $u^{\max}$ the utility derived from receiving the most beneficial program: $u^{\max} = \max\{u_k | k = 1, .., K\}$. Best and worst programs might vary among individuals. $u^{\min}$ could potentially characterize a "do nothing option", i.e. the individuals' utility without the intervention. We assume that $\mathbf{u}$ is drawn from a joint distribution $G_s(u)$ over $\mathbb{R}^K$ that depends on the value s of the sensitive attribute. We denote the random utility vector U.

An allocation policy $\mathbf{a} : \mathbb{R}^K \to \{0, 1\}^K$ assigns each individual with utility $\mathbf{u}$ to a program k if and only if $a_k(\mathbf{u}) = 1$. We assume that individuals are assigned to only one program: $\sum_{k=1}^{K} a_k(\mathbf{u}) = 1$. We denote by $\mathbf{a}.\mathbf{u}$ the inner product between the policy assignment and the individual utility: $\mathbf{a}.\mathbf{u} = \sum_{k=1}^{K} a_k(\mathbf{u})u_k$. Given N individuals i with utility $\mathbf{u}_i$, the allocation is feasible if and only if for all programs k, $\sum_{i=1}^{N} a_k(\mathbf{u}_i) \leq c_k$ (the maximum capacity for the k-th service).

3.2 Fairness, Baselines, and Normalization

In this section, we consider four notions of fairness to compare the average realized utility between groups: improvement, regret, shortfall, and gain. The definitions differ along two dimensions (1) how they normalize individual utility (additive or multiplicative), and (2) which baselines they compare individual realized utility to (worst case or best case).

The improvement and gain metrics use as a baseline the minimal or worst utility that an individual can expect from any service they receive. To be fair, the definitions say that the average increase in utility relative to the least beneficial intervention should be equal across groups. They differ in how they normalize realized utility relative to the baseline; improvement uses an additive normalization, while gain uses a multiplicative normalization.

Definition 1. ***Improvement fairness.*** *An allocation policy* $\mathbf{a}$ *satisfies fair improvement if and only if*

$$E\left[\frac{1}{N_0}\sum_{i,s=0}\mathbf{a}.(\mathbf{u}_i - u_i^{\min})\right] = E\left[\frac{1}{N_1}\sum_{i,s=1}\mathbf{a}.(\mathbf{u}_i - u_i^{\min})\right], \quad (1)$$

where the expectation is taken over samples of size N_s *for the group with sensitive attribute* $s = 0, 1$.

Definition 2. ***Gain fairness.*** *An allocation policy a satisfies fair gain if and only if*

$$E\left[\frac{1}{N_0}\sum_{i,s=0}\mathbf{a}.\frac{\mathbf{u}_i}{u_i^{\min}}\right] = E\left[\frac{1}{N_1}\sum_{i,s=1}\mathbf{a}.\frac{\mathbf{u}_i}{u_i^{\min}}\right]. \quad (2)$$

We denote by $\Delta I(\mathbf{a})$ the difference in improvement between groups:

$$\Delta I(\mathbf{a}) = E\left[\frac{1}{N_1}\sum_{i,s=1}\mathbf{a}.(\mathbf{u}_i - u_i^{\min})\right] - E\left[\frac{1}{N_0}\sum_{i,s=0}\mathbf{a}.(\mathbf{u}_i - u_i^{\min})\right]. \quad (3)$$

If $\Delta I(\mathbf{a})$ is positive, the policy $\mathbf{a}$ favors group 1; if $\Delta I(\mathbf{a})$ is negative, the policy favors group 0. We define similarly differences in gain as $\Delta G(\mathbf{a})$.

[1] We follow the convention of denoting vectors in bold type and random variables with capital letters.

Regret fairness and shortfall benchmark the realized utility in comparison to the best outcome individuals can hope for from any service (as such they are related to the classical definition of *equitability* in fair division, albeit with differences in normalization). Both fairness notions are satisfied when the average loss of utility compared to receiving the most beneficial program is equalized across groups.

Definition 3. ***Regret fairness.*** *An allocation policy a satisfies regret fairness if and only if*

$$E\left[\frac{1}{N_0}\sum_{i:s=0}\mathbf{a}.(u_i^{\max}-\mathbf{u}_i)\right]=E\left[\frac{1}{N_1}\sum_{i:s=1}\mathbf{a}.(u_i^{\max}-\mathbf{u}_i)\right], \quad (4)$$

Definition 4. ***Shortfall.*** *An allocation policy a satisfies shortfall if and only if*

$$E\left[\frac{1}{N_0}\sum_{i:s=0}\mathbf{a}.\frac{\mathbf{u}_i}{u_i^{\max}}\right]=E\left[\frac{1}{N_1}\sum_{i:s=1}\mathbf{a}.\frac{\mathbf{u}_i}{u_i^{\max}}\right]. \quad (5)$$

Like differences in improvement or in gain, we denote differences in shortfall and regret as $\Delta S(\mathbf{a})$ and $\Delta R(\mathbf{a})$, respectively. Note that $\Delta R(\mathbf{a}) \geq 0$ means that the policy $\mathbf{a}$ favors group $S=0$ over group $S=1$ for regret fairness. All four definitions represent reasonable and desirable properties of a fair allocation. However, the following results show that a decision-maker faces trade-offs when choosing which fairness notion to target. Not only might the notions not be satisfied simultaneously, it is possible to generate explicitly contradictory conclusions across the relatively similar fairness metrics regarding which group is under-served.

3.3 Improvement versus Regret

Our first result shows that improvement and regret fairness cannot be satisfied simultaneously, unless we impose strong restrictions on how groups differ. Consider two random variables $U^{\max}$ and $U^{\min}$ defined on individual most and least beneficial utility. The maximum individual utility gain that can be delivered by a service is then a random variable $\Delta U = U^{\max} - U^{\min}$.

Theorem 1. *If an allocation policy* $\mathbf{a}$ *satisfies both improvement and regret fairness then the average maximum utility gain* ΔU *must be equal across groups:* $E[\Delta U|S=0] = E[\Delta U|S=1]$. *Moreover,* $\Delta I(\mathbf{a}) + \Delta R(\mathbf{a}) = E[\Delta U|S=1] - E[\Delta U|S=0]$.

Proof. The proof is based on the following identities:

$$\begin{aligned}
\Delta I(\mathbf{a}) &= E\left[\frac{1}{N_1}\sum_{i,s=1}\mathbf{a}(\mathbf{u}).(\mathbf{u}_i - u_i^{\max}+\Delta u_i)\right] \\
&- E\left[\frac{1}{N_0}\sum_{i,s=0}\mathbf{a}(\mathbf{u}).(\mathbf{u}_i - u_i^{\max}+\Delta u_i)\right] \\
&= E\left[\frac{1}{N_1}\sum_{i,s=1}\sum_{k=1}^{K}\mathbf{a}_k(\mathbf{u})\Delta u_i\right] \quad (6)\\
&- E\left[\frac{1}{N_0}\sum_{i,s=0}\sum_{k=1}^{K}\mathbf{a}_k(\mathbf{u})\Delta u_i\right] - \Delta R(\mathbf{a}) \\
&= E\left[\frac{1}{N_1}\sum_{i,s=1}\Delta u_i\right] - E\left[\frac{1}{N_0}\sum_{i,s=0}\Delta u_i\right] - \Delta R(\mathbf{a}),
\end{aligned}$$

where the last equality comes from the fact that $\sum_{k=1}^{K} a_k(u) = 1$ for all u. Therefore, if $\Delta I(\mathbf{a}) = 0$ and $\Delta R(\mathbf{a}) = 0$, then $E[\Delta U|S=0] = E[\Delta U|S=1]$. □

The result in Theorem 1 implies that regardless of the allocation policy, for both improvement and regret fairness to hold it is necessary that groups would gain on average similarly if they were always allocated their most beneficial intervention. Thus, a trade-off exists when defining what a fair assignment should look like: for example, a policy satisfying improvement fairness would always violate regret fairness unless $E[\Delta U|S=0] = E[\Delta U|S=1]$. Since $\Delta I(\mathbf{a}) + \Delta R(\mathbf{a}) = E[\Delta U|S=1] - E[\Delta U|S=0]$, the closer a policy is to satisfying improvement fairness, the worse its regret fairness, and vice-versa. A follow up question is whether improvement and regret fairness tell a different story about the fairness of an allocation policy a. The next result shows that whenever $E[\Delta U|S=0]$ and $E[\Delta U|S=1]$ differ, unless all policies favor one group, there exists a policy that favors one group for improvement fairness and favors the other one for regret fairness.

Theorem 2. *Suppose that* $E[\Delta U|S=1] > E[\Delta U|S=0]$. *Suppose that there exists a policy that favors group* $S=0$ *for improvement fairness and another policy that favors group* $S=1$ *for improvement fairness. Then, there exists a policy* $\mathbf{a}^*$ *such that* $\Delta I(\mathbf{a}^*) > 0$ *and* $\Delta R(\mathbf{a}^*) > 0$. *That is, there exists a policy that favors* $S=1$ *with respect to improvement fairness (larger is better), but favors* $S=0$ *with respect to regret fairness (lower is better).*

The proof of Theorem 2 relies on the fact that the set of differences in improvement/regret is a continuous interval:

Lemma 1. *Suppose that there exist two allocation policies* $\mathbf{a}$ *and* $\mathbf{a}'$ *with differences in improvement* δ *and* $\delta' > \delta$. *Then, for any* $\delta^* \in [\delta, \delta']$, *there exists an allocation policy* $\mathbf{a}^*$ *with difference in improvement equal to* δ^*. *A similar result holds for difference in regret.*

Proof. We show the result for differences in improvement. The proof can be readily extended to differences in regret. We choose $\lambda = \frac{\delta'-\delta^*}{\delta'-\delta} \in [0,1]$. We define an allocation policy $\mathbf{a}^\lambda$ as follows:

- Partition randomly the individuals into two populations G_λ and $G_{1-\lambda}$ of size λN and $(1-\lambda)N$, respectively.
- For each program k, assign λc_k of them to the population G_λ; and $(1-\lambda)c_k$ of them to the population $G_{1-\lambda}$.
- Apply the allocation policy $\mathbf{a}$ to the population G_λ and $\mathbf{a}'$ to the population $G_{1-\lambda}$.

By construction the policy a_λ satisfies the resource constraints. Moreover,

$$\Delta I(\mathbf{a}_\lambda) = \Delta I(\mathbf{a})P(G_\lambda) + \Delta I(\mathbf{a}')(1-P(G_{1-\lambda}) = \delta\lambda + \delta'(1-\lambda) = \delta^*, \quad (7)$$

where the last equality comes from our choice for the value of λ. □

Proof. Theorem 2. We choose $\epsilon = \frac{E[\Delta U|S=1]-E[\Delta U|S=0]}{2}$. $\epsilon > 0$ by assumption. Using the assumption of Theorem 2, there exist $\mathbf{a}$ and $\mathbf{a}'$ such that $\Delta I(\mathbf{a}) < 0$ and $\Delta I(\mathbf{a}') > 0$. We apply Lemma 1 with $\delta = \Delta I(\mathbf{a}) < 0$, $\delta' = \Delta I(\mathbf{a}') > 0$ and $\delta^* = \min\{\epsilon, \delta'/2\}$: there exists a policy $\mathbf{a}^*$ such that $\Delta I(\mathbf{a}^*) = \delta^* > 0$. Moreover, $\Delta R(\mathbf{a}^*) = E[\Delta U|S=1] - E[\Delta U|S=0] - \Delta I(\mathbf{a}^*) \geq \epsilon > 0$. □

Thus, regret fairness and improvement fairness cannot hold simultaneously unless populations are homogeneous in terms of their best response to the allocation (Theorem 1). Moreover, assessing which group is favored by a given policy can lead to contradictory results depending on whether we measure the fairness properties of the policy in terms of differences in improvement or regret. The result in Theorem 2 illustrates that decision-makers cannot expect that improvement and regret notions tell a similar story about whether an allocation policy under-serves a given group. Results Theorem 1 and Theorem 2 are general, since they hold for any set of capacities c_1, ..., c_K and for distributions of utilities provided that $E[\Delta U|S=1] > E[\Delta U|S=0]$. Both illustrate the central role of the difference between $E[\Delta U|S=0]$ and $E[\Delta U|S=1]$ in driving wedges between improvement and regret fairness. Additionally, Theorem 2 is not very restrictive in its assumptions, since it only requires that no group is under-served regardless of the policy.

3.4 Shortfall versus Gain

In this section, we show that the fairness trade-offs between improvement and regret exist also with multiplicative notion of fairness, gain and shortfall. Unlike trade-offs between improvement and regret where our results are general, in the case of shortfall versus gain, we derive results in a stylized framework and leave it to future work to extend our results to more general settings. Nevertheless, this section captures the essence of the problem in the multiplicative setting. We denote for each individual by $r = u^{\min}/u^{\max}$ the ratio between the lowest and highest utility obtained from the intervention. This serves as a multiplicative counterpart of Δu. We consider the following framework (SF1):

- There are two types of individuals: type A with high value $\overline{r}$ for the ratio r; type B with a low value $\underline{r} < \overline{r}$ for r.
- Conditional on r, the distribution of utility is similar across programs and types.

In this stylized framework, assigning to an individual their most beneficial program delivers either a large increase $\overline{r}$ over $u^{\min}$ (type A) or a lower one $\underline{r}$ (type B). We characterize the heterogeneity across groups by differences in the distribution of type A and B within each group. We denote by π_0 the proportion of type B individuals for group $S=0$; and, π_1 the proportion of type B for group $S=1$.

Theorem 3. *In the stylized framework (SF1):*

- *A policy satisfies both shortfall and gain fairness if and only if $\pi_0 = \pi_1$.*
- *If $\pi_0 \neq \pi_1$, a policy a that achieves gain (shortfall) fairness, favors, according to shortfall (gain) fairness whichever group has the lowest proportion of type A individuals.*

Proof. In this proof, let $\underline{\alpha}$ denote $E\left[\frac{\mathbf{a}(u).\mathbf{u}}{u^{\min}}|r=\underline{r}\right]$ and $\overline{\alpha}$ denote $E\left[\frac{\mathbf{a}(u).\mathbf{u}}{u^{\min}}|r=\overline{r}\right]$. Then, we write (for any policy) differences in gain,

$$\Delta G(\mathbf{a}) = \left\{\pi_1\underline{\alpha} + (1-\pi_1)\overline{\alpha}\right\} - \left\{\pi_0\underline{\alpha} + (1-\pi_0)\overline{\alpha}\right\} = (\pi_0-\pi_1)(\overline{\alpha}-\underline{\alpha}) \tag{8}$$

and differences in shortfall as

$$\begin{aligned}\Delta S(\mathbf{a}) &= \left\{\pi_1\underline{\alpha r} + (1-\pi_1)\overline{\alpha r}\right\} - \left\{\pi_0\underline{\alpha r} + (1-\pi_0)\overline{\alpha r}\right\} \\ &= (\pi_0-\pi_1)(\overline{\alpha r} - \underline{\alpha r}).\end{aligned} \tag{9}$$

Therefore, gain and shortfall fairness are equivalent to $(\pi_0-\pi_1)(\overline{\alpha}-\underline{\alpha}) = 0$ and $(\pi_0-\pi_1)(\overline{\alpha r}-\underline{\alpha r}) = 0$. Hence, if $\pi_0 \neq \pi_1$, $\overline{\alpha} = \underline{\alpha}$ and $\overline{\alpha}\,\overline{r} = \underline{\alpha}\,\underline{r}$, which is not possible since $\underline{r} \neq \overline{r}$.

To show the second part of Theorem 3, we use the fact that gain fairness implies that $\overline{\alpha} = \underline{\alpha}$ (equation (8)) and that the difference in shortfall between group $S=1$ and $S=0$ can be then written $\Delta S(\mathbf{a}) = (\pi_0-\pi_1)(\overline{r}-\underline{r})\overline{\alpha}$, which have the same sign as $\pi_0-\pi_1$ since $\overline{r} > \underline{r}$. Therefore, if $\pi_0 > \pi_1$, the policy favors group $S=1$ with respect to shortfall fairness; otherwise, it favors group $S=0$. □

Theorem 3 states that shortfall and gain can be satisfied simultaneously if and only if populations have similar fractions of type A individuals. It is similar in spirit to the results above, showing that unless populations meet stringent requirements of similarity in utility distributions between groups (in this case instantiated by the fractions of the two types in each population), the versions of fairness characterized by comparing with the min versus the max cannot be simultaneously satisfied.

3.5 Multiplicative versus Additive Normalization

Improvement and gain fairness aim at capturing a similar fairness concept: groups receive on average the same increase in utility relative to assigning the least beneficial service. Both fairness metrics differ only by whether the normalization relative to the lowest utility that an individual can derive from the overall intervention is additive or multiplicative. In this section, we show that even the choice of normalization generates inherent fairness trade-offs.

We consider the following stylized framework (SF2):

- There are two types of individuals: type C for which $u^{\min}$ takes a low value $\underline{u}$; and type D for which $u^{\min}$ takes a larger value $\overline{u} > \underline{u}$.
- Conditional on $u^{\min}$, the distribution of utility is similar across programs and types.

Although stylized, both assumptions allow us to characterize the heterogeneity across groups by differences in their distribution over $u^{\min}$. Let p_s denote the fraction of type C for group $S=s$. Differences in p_s across groups imply differences in the distribution of utility $P(U|S)$ within each group, even if the conditional distribution $P(U|U^{\min})$ is similar across types.

Theorem 4. *In the stylized framework (SF2) with types C and D, a policy* **a** *satisfies both improvement fairness and gain fairness for group $S=0$ and $S=1$ if and only if one of the following conditions holds:*

- *$p_0 = p_1$;*
- *the policy* **a** *assigns the least beneficial program to everyone (i.e. $\mathbf{a}.\mathbf{u} = u^{\min}$).*

Proof. Let $\underline{\beta}$ denote $E[\mathbf{a}.\mathbf{u}|U^{min} = \underline{u}]$ and $\overline{\beta}$ denote $E[\mathbf{a}.\mathbf{u}|U^{min} = \overline{u}]$. Then, we write differences in improvement as

$$\begin{aligned}\Delta I(\mathbf{a}) &= \left\{p_1\underline{\beta} + (1-p_1)\overline{\beta} - p_1\underline{u} - (1-p_1)\overline{u}\right\} \\ &\quad - \left\{p_0\underline{\beta} + (1-p_0)\overline{\beta} - p_0\underline{u} - (1-p_0)\overline{u}\right\} \\ &= (p_1-p_0)(\underline{\beta} - \overline{\beta} + \underline{u} - \underline{u})\end{aligned} \tag{10}$$

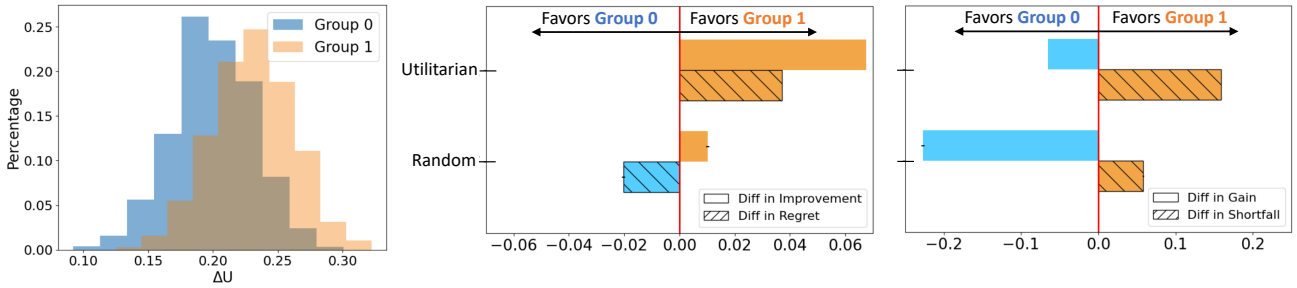

(a) Distribution of ΔU (b) Improvement vs. Regret (c) Gain vs. Shortfall

Figure 1: Simulation results when groups have different mean of utilities. Panel (a) shows the distribution of the maximum utility gains $\Delta U = U^{\max} - U^{\min}$ for *group 0* (blue), and *group 1* (orange). Panel (b) shows the differences in improvement and regret, Panel (c) shows the differences in gain and shortfall. Error bars show the 95% confidence interval of each fairness metric over 100 instantiations of the random allocation.

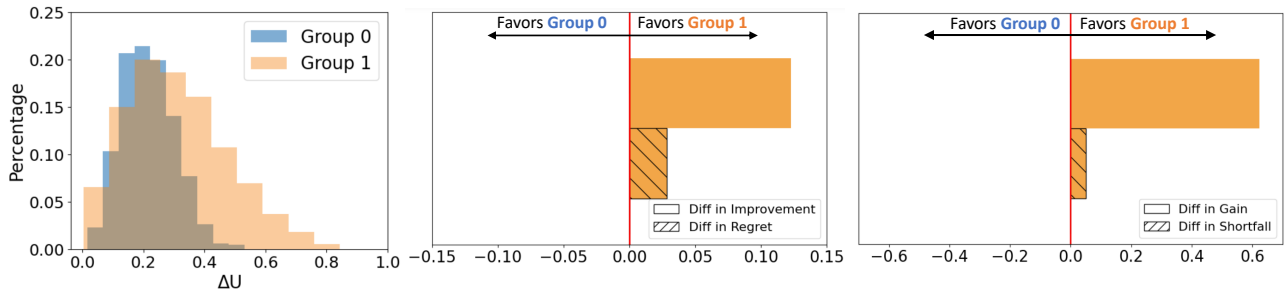

(a) Distribution of ΔU (b) Improvement vs. Regret (Utilitarian) (c) Gain vs. Shortfall (Utilitarian)

Figure 2: Simulation results when groups have the same mean utilities for the services, but different variances. Panel (a) shows the distribution of the maximum utility gains $\Delta U = U^{\max} - U^{\min}$ for *group 0* (blue), and *group 1* (orange). Panel (b) shows the differences in improvement and regret, Panel (c) shows the differences in gain and shortfall. Group 1 is favored strongly by all the fairness measures when allocations are utilitarian.

and differences in gain as

$$\begin{aligned}\Delta G(\mathbf{a}) &= \left\{p_1 \frac{\underline{\beta}}{\underline{u}} + (1-p_1)\frac{\overline{\beta}}{\overline{u}}\right\} - \left\{p_0 \frac{\underline{\beta}}{\underline{u}} + (1-p_0)\frac{\overline{\beta}}{\overline{u}}\right\} \\ &= (p_1 - p_0)\left(\frac{\underline{\beta}}{\underline{u}} - \frac{\overline{\beta}}{\underline{u}}\right)\end{aligned} \tag{11}$$

Therefore, improvement fairness and gain are equivalent to $(p_0 - p_1)(\underline{\beta} - \overline{\beta} + \underline{u} - \underline{u}) = 0$ and $(p_0 - p_1)\left(\frac{\underline{\beta}}{\underline{u}} - \frac{\overline{\beta}}{\underline{u}}\right) = 0$. If $p_0 \neq p_1$, improvement and gain fairness imply $\underline{\beta} = \frac{\underline{u}}{\overline{u}}\overline{\beta}$ and $\overline{\beta} = \overline{u}$. The latter equality leads to $\mathbf{a}.\mathbf{u} = \overline{u}$ if $u^{\min} = \overline{u}$ and the former equality leads to $\mathbf{a}.\mathbf{u} = \underline{u}$ if $u^{\min} = \underline{u}$. □

Theorem 4 demonstrates a simple, yet general, setting where improvement fairness and gain fairness cannot be obtained simultaneously unless either the distribution of utilities are the same across groups ($p_0 = p_1$) or the policy does not create any utility improvement relative to $U^{\min}$.

4 SIMULATIONS WITH UTILITARIAN AND RANDOM ALLOCATIONS

Thus far, we have not needed to define an allocation policy explicitly, since we were focused on existence results. In this section, we consider two natural allocation policies – utilitarian (maximizing the sum of utilities of all agents) and random. Both must respect capacity constraints. We simulate a simple environment with two groups and three services. In one setting, members of the two groups have different mean utilities from receiving the three services, while the variances are the same. In the second, members of the two groups have the same mean utilities from receiving the three services, but different variances. We are interested in understanding (1) how the different fairness measures behave in these two settings; (2) the role played by utilitarian objectives in the assignment problem.

In our setting, there are three ($k = 1, 2, 3$) services with fixed capacities ($c_1 = c_2 = c_3 = 1000$) and 3000 applicants divided into two groups of equal size: *group 0* and *group 1*. We sample individual utilities for service k from a normal distribution with mean μ_{sk} and standard deviation σ_{sk}, $s = 0$ for *group 0* and $s = 1$ for *group 1*.

4.1 Groups with Different Means

In this set of simulations, we study the behavior of fairness measures when individual utilities are sampled from group-dependent distributions. The groups have different sample means μ but the same variances σ^2. For *group 0*, the means of the three services are $\mu_{01} = 0.2, \mu_{02} = 0.3,$ and, $\mu_{03} = 0.4$. For *group 1*, the means are $\mu_{11} = 0.4, \mu_{12} = 0.5,$ and, $\mu_{13} = 0.63$ The variances of the three services for both groups are equal, $\sigma^2_{01} = \sigma^2_{11} = 1 \times 10^{-4}$, $\sigma^2_{02} = \sigma^2_{12} = 4 \times 10^{-4}$, and, $\sigma^2_{03} = \sigma^2_{13} = 9 \times 10^{-4}$. Individuals in *group 1* have on average higher utilities for all services.

As pointed out in section 3.3, we observe in Figure 1 that the difference in ΔU leads to a trade-off between the improvement and regret fairness metrics. Figure 1 shows that even for a random assignment, different metrics lead to conflicting fairness assessment. The improvement fairness metric favors the group with higher mean ΔU (*group 1*), and regret favors the groups with lower mean ΔU (*group 0*). To complicate fairness assessment further, switching from additive to multiplicative normalization reverses which group is favored.

Moreover, the utilitarian allocation appears to favor *group 1* according to improvement, regret and gain, but favors *group 0* in terms of shortfall. These results confirm in a simulated environment that utility normalization has profound implications on how we assess the fairness of an allocation.

4.2 Groups with Equal Means and Different Variances

In our second set of simulations, we study the effects of groups having similar means but different variances, a situation that is commonly discussed, for instance in the context of gender differences in student performance [7]. In this case, we hypothesize that the higher variance group is likely to be favored by utilitarian allocations. For both groups, the means for the three services are equal, $\mu_{01} = \mu_{11} = 0.4$, $\mu_{02} = \mu_{12} = 0.5$, and $\mu_{03} = \mu_{13} = 0.6$. For *group 0*, the variances for the three interventions are set to $\sigma^2_{01} = 9 \times 10^{-5}$, $\sigma^2_{02} = 2 \times 10^{-3}$, $\sigma^2_{03} = 1 \times 10^{-2}$, while for *group 1*, the variances for the three interventions are set to $\sigma^2_{11} = 9 \times 10^{-3}$, $\sigma^2_{12} = 2 \times 10^{-2}$, $\sigma^2_{13} = 3 \times 10^{-2}$. Thus, *group 0* has lower variance.

Our results in Figure 2 show that, as hypothesized, the group with larger variance (group 1) is indeed favored according to all fairness metrics. When maximizing the sum of utilities, it is optimal to assign their best services to individuals with utilities in the tail of the distribution. We find that a larger fraction (65%) of individuals in *group 1* than in *group 0* (46%) receive the service that maximizes their utility.

5 FAIRNESS TRADE-OFFS IN HOMELESS SERVICE DELIVERY

Our theoretical analysis suggests that heterogeneity in service responses across groups drives fairness metrics in opposite directions. In this section, we investigate whether the fairness tradeoffs emerge in the capacitated assignment of homeless services across several sub-populations. We hypothesize that if sociodemographic group differences exist in the utilities received from allocations (and in particular, between the differences in the best versus worst allocations), then we should see tradeoffs between improvement versus regret fairness, shortfall versus gain, and improvement versus gain. We provide evidence for both the antecedent (*heterogeneity in responses across groups*) and the consequent (*inherent fairness trade-offs between groups*).

5.1 Background

Homelessness represents a socioeconomic and public health challenge for many communities in the United States. Approximately 1.5 million people experience homelessness for at least one night every year [23, 30]. Homelessness has short- and longer-term implications on health, employment, and crime [14, 26, 34]. Guided by federal policies, communities offer an array of services for households lacking stable and permanent living accommodations. We study three main homeless services: Transitional Housing (TH); Rapid Rehousing (RRH) and Emergency Shelter (ES). Transitional Housing provides accommodation for up to 24 months with comprehensive case management to address barriers toward stable housing, such as substance abuse and issues related to behavioral health. Rapid Rehousing offers access to rental units for six months without intensive case management. Emergency Shelter provides a bed to sleep at night for no more than one or two months. On a daily basis, caseworkers assign homeless households seeking assistance to an available service, reserving the most intensive TH for those with greater needs.

5.2 Data

Our main dataset is based on estimated probabilities of households re-entering homelessness services within two years after initial receipt of services. This data, collected by [36] is publicly available.[2] The estimates are based on applying a machine learning model (BART [31]) to administrative records that tracked service provision in a metropolitan area from 2007 through 2014. Service providers collected demographic and household characteristics upon entry into the system, and data capture the intervention assigned and whether households subsequently requested additional assistance [36]. The model estimates counterfactual probabilities p_{ik} of a household i to re-enter the homeless system within 2 years given the assignment of a specific service k, where $k \in \{TH, RRH, ES\}$. The original data also tracks responses to homelessness prevention – time-limited monetary assistance that differs from the other three interventions that allocate actual bed space. Given that the constraints on homelessness prevention are different, we focus here only on households that needed actual bed space (and were therefore not eligible to receive prevention services). Therefore, our final data contains 3, 375 households and they received either TH, RRH, or ES.

We compute the utility of service k to individual i as $u_{ik} = 1 - p_{ik}$. We obtained from Kube et al additional sociodemographic characteristics for each household, including race, gender, age, disability status, presence of spouse and/or children, and household size. We note that our analysis takes the counterfactual probability estimates as given. These estimates are based on an ignorability assumption implicit in using BART in this manner [31, 36].

[2] https://github.com/amandakube/Allocating-Homelessness-Interventions---Counterfactual-Predictions

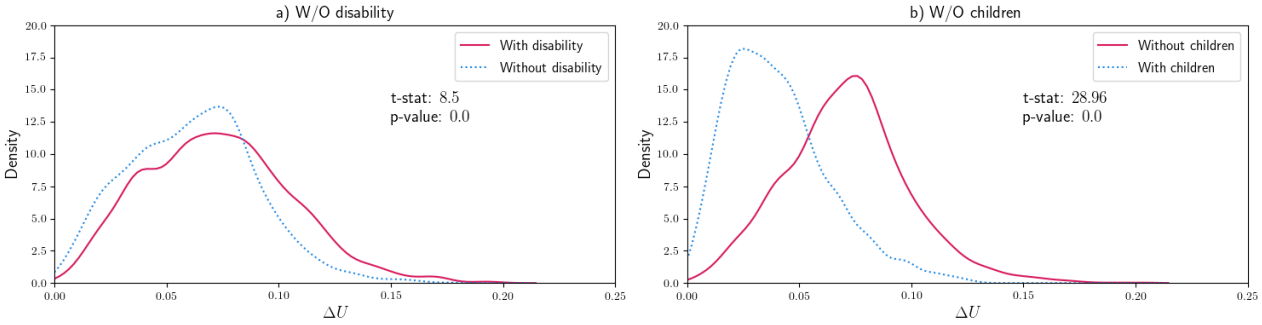

Figure 3: Distribution of the maximum utility gain ΔU that individuals can derive from the homeless system across various demographic groups. We obtain the probability density function of $\Delta U = U^{max} - U^{min}$ via Gaussian kernel density estimation with a bandwidth of 0.2. Differences in probability density functions between households with and without disability (Panel a)) and with and without spouse (Panel b)) illustrate heterogeneous responses to housing assistance.

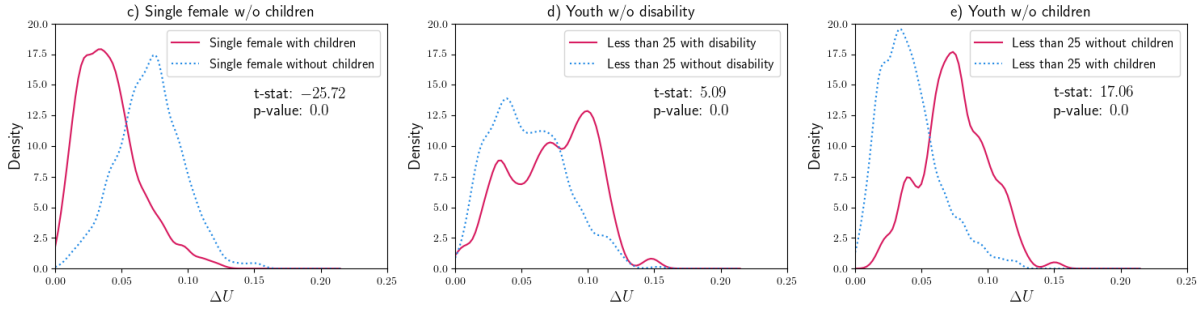

Figure 4: Same as Figure 3 but for intersection groups single female with and without children (Panel c)); youth under 25 with and without disability (Panel d)); and, youth under 25 with and without children (Panel e)).

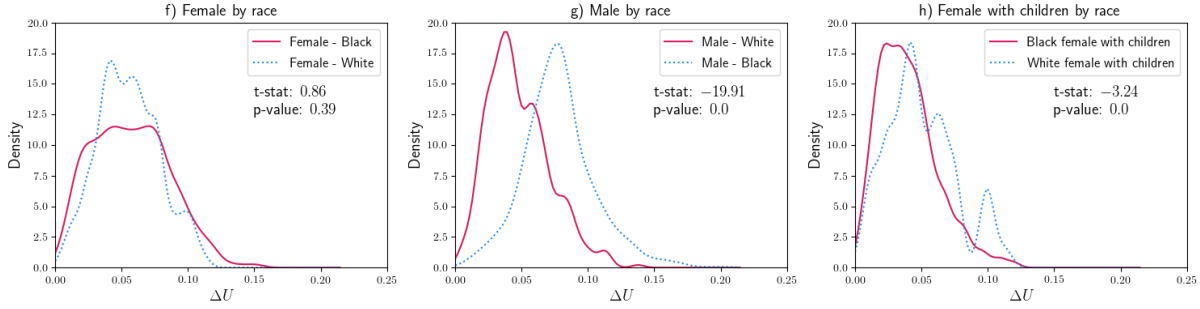

Figure 5: Same as Figure 3 but for groups defined by perceived racial background.

We define a series of sociodemographic groups and intersectional identities expected to exhibit substantial heterogeneity in responses to homeless services. First, households with disabilities are considered more vulnerable, and prior research shows that more vulnerable households do best with more intensive services [4, 40]. Therefore, we expect households with disabilities to benefit more from TH and less from ES. Second, families with children under the age of 18 experience homelessness due to socioeconomic reasons rather than disability and vulnerability, and thus, we anticipate families will respond better to rapid rehousing than more intensive TH [16, 22, 43]. Third, we examine the intersection between gender and family status, assuming that single female households without children do better in TH compared with single female-headed families with children, who are more likely to benefit from RRH. Fourth, we look within households headed by youth aged 18 to 24 years to compare disability status (versus no disability) and family status (children versus no children), hypothesizing that those with disabilities benefit more from TH and families with children from

Table 1: Distribution of services that deliver to each household the highest utility across demographic groups. This shows the fraction of households in each demographic group for which ES, TH or RRH leads to the lowest probability to re-enter the homeless system.

	TH	RRH	ES	Size		TH	RRH	ES	Size
All	0.68	0.27	0.05	3375					
With disability	0.73	0.23	0.03	601	Without disability	0.66	0.28	0.06	2699
Without children	0.85	0.14	0.01	2533	With children	0.16	0.67	0.17	842
Single female with children	0.15	0.7	0.15	761	Single female without children	0.7	0.3	0.01	885
Less than 25 with disability	0.62	0.37	0.01	78	Less than 25 without disability	0.47	0.49	0.05	501
Less than 25 without children	0.83	0.17	0.0	275	Less than 25 with children	0.19	0.73	0.08	308
Female - Black	0.46	0.46	0.08	1414	Female - White	0.34	0.62	0.04	241
Male - Black	0.95	0.02	0.03	1254	Male - White	0.8	0.14	0.06	406

RRH [39]. Lastly, given the over-representation in homelessness of minorities and especially Black households, we test how race affects homeless service utilities [30].

Prior research suggests the causes of homelessness vary for White people experiencing disabilities, versus Black people experiencing greater housing discrimination [33]. Moreover, race intersects with gender (males vs females) and family status (with children versus without children) in ways that could drive variation in homeless service outcomes.

5.3 Heterogeneity across Demographic Groups

In this section, we document heterogeneity in the distributions of utility across various groups. For each household, we compute the difference ΔU between its best and worst utility. Figure 3 shows heterogeneity in response to homeless services across households with and without reported disabilities; with and without children. The distribution of ΔU for households with a disability skews to the right (panel a); assigning the best service to a disabled client has a larger impact in terms of re-entry probability than assigning a client without a disability to their most beneficial service. The difference in distribution means is statistically significant with a t-statistic of 8.5 and p-value infinitesimally small. This finding aligns with prior research that shows vulnerable households do best with more intensive services [4, 40]. The distribution of ΔU for households without children skews strongly to the right compared with households with children (panel b). The mean of ΔU for households without children is 0.07, while it is only 0.04 for household with children. The difference is statistically significant with a t-statistic of 29.0 and a p-value infinitesimally small. This result illustrates how families with children differ in their responses to housing assistance compared to homeless individuals.

Figure 4 looks at intersectional sociodemographic groups. Panel c shows that the impact of different homeless services for a single female depends strongly on whether there are children in the household. Similarly, youth with and without disability respond differently to homeless services (panel d). For both intersections, the difference in means is statistically significant (t-statistic of 25.7 for single female vs single mother and 5.1 for youth with vs without disability).

Figure 5 explores differential responses to housing assistance by race and shows substantial differences in the distribution of ΔU between Black and White males (Panel g). Black homeless populations may on average benefit more from more intensive homeless services. Prior research [33] suggests that social discrimination and socio-economic disadvantage could increase the risk for homelessness among populations with perceived Black background and that housing assistance could mitigate some of these vulnerabilities.

Results from Figures 3, 4 and 5 suggest that heterogeneity in utility is pervasive. Table 1 explains some of this heterogeneity by identifying which of the three services (TH, RRH and ES) benefits the most households within each group. For the homeless population studied in this paper, TH is the most preferred service for 68% of the population, followed by RRH (27%) and ES (5%). This preference for more intensive care is exacerbated for households with disability (73% prefer TH), in line with prior findings that most vulnerable populations benefit from more integrated care. The preferences of households with a disability toward TH contrasts with the preferences of families with children toward RRH: 67% of households with children benefit most from RRH, while TH is best for only 16% of families. This holds true for all intersectional groups that include children and could explain differences between males and females, since females are more likely to live with children. Regardless of gender, TH is likely most beneficial for the Black homeless population: TH is the most beneficial service for 46% of Black females but only for 34% of White females; and, for 95% of Black males but only for 80% of White males.

5.4 Fairness Trade-Offs in the Observed Allocation of Homeless Services

Our theory suggests that heterogeneity in the distribution of the maximum gain ΔU would drive fairness metrics in opposite directions: (i) there exist assignments of homeless services with conflicting fairness assessment depending on choosing improvement, regret, gain or shortfall as the fairness metric (Theorem 2); (ii) assignments that satisfy improvement fairness could violate regret fairness and vice-versa (Theorem 1). Given the substantial heterogeneity among the sociodemographic and intersectional groups

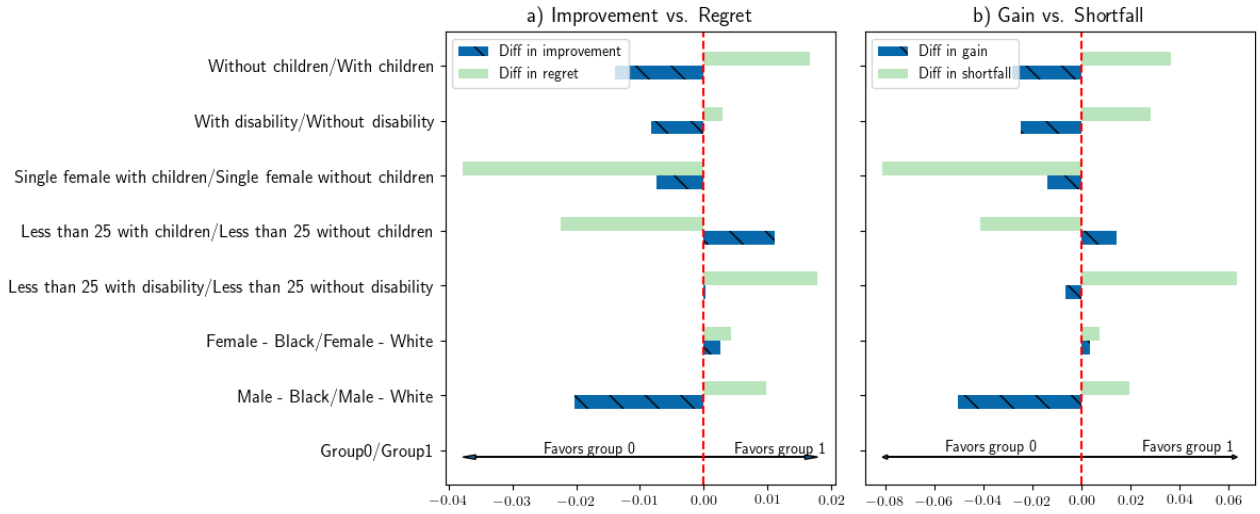

Figure 6: Fairness trade-off in the observed assignment of homeless services. This compares which demographic group is favored by the assignment depending on the fairness metric. Trade-offs occur when improvement favors one group and regret the other one (left panel) or when shortfall favors one group and gain the other (right panel).

documented in section 5.3, Theorem 2 implies that ambiguous fairness assessments can arise for some policies. However, Theorem 2 is not constructive. Are such policies realistic in the context of homeless services delivery? Here we test whether the observed assignment as reported in the administrative records is subject to contradictory fairness assessments depending on the choice of fairness metric.

Figure 6 (Panel a) plots the difference in improvement ΔI and the negative of difference in regret $-\Delta R$. Positive values indicate that the policy favors group $S = 1$, while negative values mean the policy favors group $S = 0$. According to improvement, the observed assignment favors households without children, while according to regret, it favors households with children: $\Delta I = -0.013$, while $-\Delta R = 0.016$. A similar ambiguity emerges for households with and without disability. Moreover, choosing improvement or regret flips the conclusion on whether the assignment is unfair to Black males relative to White males: Black males derive higher utility gains according to improvement ($\Delta I = -0.02$) but lower according to regret ($-\Delta R = 0.009$). The results provide empirical evidence that policies that lead to contradictory fairness assessment are not just theoretical oddities, but do occur in real world applications. Empirically, similar trade-offs occur for shortfall versus gain (Figure 6, Panel b). Moreover, in Figure 6, we find one pairwise comparison, youth with versus without a disability, for which the observed policy satisfies improvement fairness. This instance of improvement fairness allows us to test whether Theorem 1 holds here. We find that the policy does not satisfy regret fairness, consistent with the heterogeneity in ΔU found in section 5.3 between youth with versus without a disability.

6 CONCLUSION

How do we judge whether an approach to allocation of scarce societal resources is fair for different sociodemographic groups of public concern? The problem lies at the intersection of recent work in fair machine learning and a long history of work from economics, social choice, and algorithmic game theory on fair division. It also brings into question concerns of local justice [20], which studies how individuals are prioritized in the allocation of scarce resources by local institutions. The key point we make in this paper is that *baselines matter when we measure outcomes for different groups*. The exact same allocation may favor one group over another when assessed against the baseline intervention of doing nothing, but the group it favors could invert when measured against the baseline of giving each group the best intervention it could get in a scenario with no resource constraints. The social objective being optimized also can drive fairness results – for example, utilitarian allocations typically favor groups with higher variance in utilities across different types of services, even if the means are the same.

Our results are more than theoretical. We show that the pattern arises in homeless service delivery, where outcomes vary by and within sociodemographic groups. For instance, returns to homelessness vary by service allocation more for households without children compared to families with children. Naive policy applications that fail to consider baseline variation may negatively impact some groups. Aiming to reduce overall homelessness, for example, by prioritizing households without children for intensive services disproportionately excludes households with children from receiving their best service, whereas an alternative policy that matches households with children to their best service fails to reduce overall homelessness. The data illustrate similar fairness tradeoffs across intersecting sociodemographic groups, including disability status, gender, age, and race. Failing to consider carefully the underlying distributions and metrics for success threatens counterproductive policy initiatives. Current national advocacy to reduce veteran and chronic homelessness to zero ask communities to shift resources in ways that may undermine other goals [2]. Moreover, federal and

local policies simultaneously strive for system efficiency and equity, which prove antithetical in many contexts [26]. Our findings raise serious questions for institutions when designing homeless policies and social policy more generally.

ACKNOWLEDGMENTS

We thank Nisarg Shah for helpful conversations. We are grateful for support from the NSF through awards 2127752, 2127754, and 1939677, and from Amazon through an NSF-Amazon Fairness in AI award. Special thanks go to the homeless service consumers and providers represented in the data.

REFERENCES

[1] [n.d.]. HUD System performance measures. https://www.hudexchange.info/programs/coc/system-performance-measures/#guidance Published online at https://www.hudexchange.info/programs/coc/system-performance-measures/#guidance.
[2] 2021. Functional zero. https://community.solutions/built-for-zero/functional-zero/ By Community Solutions, available at https://community.solutions/built-for-zero/functional-zero/.
[3] Kenneth J Arrow. 1950. A difficulty in the concept of social welfare. *Journal of Political Economy* 58, 4 (1950), 328–346.
[4] Tim Aubry, Gary Bloch, Vanessa Brcic, Ammar Saad, Olivia Magwood, Tasnim Abdalla, Qasem Alkhateeb, Edward Xie, Christine Mathew, Terry Hannigan, Chris Costello, Kednapa Thavorn, Vicky Stergiopoulos, Peter Tugwell, and Kevin Pottie. 2020. Effectiveness of permanent supportive housing and income assistance interventions for homeless individuals in high-income countries: A systematic review. *The Lancet Public Health* 5 (2020), e342–e360. https://doi.org/10.1016/S2468-2667(20)30055-4
[5] Haris Aziz. 2020. Justifications of welfare guarantees under normalized utilities. *ACM SIGecom Exchanges* 17, 2 (2020), 71–75.
[6] John J Bartholdi III, Craig A Tovey, and Michael A Trick. 1992. How hard is it to control an election? *Mathematical and Computer Modelling* 16, 8-9 (1992), 27–40.
[7] Ariane Baye and Christian Monseur. 2016. Gender differences in variability and extreme scores in an international context. *Large-scale Assessments in Education* 4, 1 (2016), 1–16.
[8] Sylvain Bouveret and Michel Lemaître. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems* 30, 2 (2016), 259–290.
[9] Molly Brown, Camilla Cummings, Jennifer Lyons, Andrés Carrión, and Dennis P Watson. 2018. Reliability and validity of the Vulnerability Index-Service Prioritization Decision Assistance Tool (VI-SPDAT) in real-world implementation. *Journal of Social Distress and the Homeless* 27, 2 (2018), 110–117.
[10] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–32.
[11] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* 5, 2 (2017), 153–163.
[12] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*. PMLR, 134–148.
[13] Alexandra Chouldechova and Aaron Roth. 2018. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810* (2018).
[14] Elior Cohen. 2020. *Housing the Homeless: The Effect of Housing Assistance on Recidivism to Homelessness, Economic, and Social Outcomes*. Technical Report. mimeo.
[15] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. 2019. Group fairness for the allocation of indivisible goods. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1853–1860.
[16] Mary Cunningham, Sarah Gillespie, and Jacqueline Anderson. 2015. Rapid rehousing. *Urban Institute: Washington, DC, USA* (2015).
[17] Sanmay Das. 2022. Local Justice and the Algorithmic Allocation of Scarce Societal Resources. *Proceedings of the AAAI Conference on Artificial Intelligence* (2022). To appear.
[18] David Delacrétaz, Scott Duke Kominers, Alexander Teytelboym, et al. 2019. *Matching mechanisms for refugee resettlement*. Technical Report. Working paper.
[19] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. 214–226.
[20] Jon Elster. 1992. *Local justice: How institutions allocate scarce goods and necessary burdens*. Russell Sage Foundation.
[21] Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. Runaway feedback loops in predictive policing. In *Conference on Fairness, Accountability and Transparency*. PMLR, 160–171.
[22] Angela R Fertig and David A Reingold. 2008. Homelessness among at-risk families with children in twenty American cities. *Social Service Review* 82, 3 (2008), 485–510.
[23] Cassie Fisher, Brandon Sorenson Gray, Karen Quackenbush, Tamera Kohler, and Jonathan Hardy. 2018. Homelessness in the United States. *Development* (2018), 14.
[24] Marc Fleurbaey, Kotaro Suzumura, and Koichi Tadenuma. 2005. The informational basis of the theory of fair allocation. *Social Choice and Welfare* 24, 2 (2005), 311–341.
[25] Duncan Karl Foley. 1966. *Resource allocation and the public sector*. Yale University.
[26] Patrick J Fowler, Peter S Hovmand, Katherine E Marcal, and Sanmay Das. 2019. Solving homelessness from a complex systems perspective: insights for prevention responses. *Annual Review of Public Health* 40 (2019), 465–486.
[27] Josh Gardner, Christopher Brooks, and Ryan Baker. 2019. Evaluating the Fairness of Predictive Student Models Through Slicing Analysis. In *Proceedings of the 9th International Conference on Learning Analytics & Knowledge*. ACM, 225–234.
[28] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems* 29 (2016), 3315–3323.
[29] Hoda Heidari, Michele Loi, Krishna P Gummadi, and Andreas Krause. 2019. A moral framework for understanding fair ML through economic models of equality of opportunity. In *Proceedings of the Conference on Fairness, Accountability and Transparency*. 181–190.
[30] Meghan Henry, Anna Mahathey, and Meghan Takashima. 2020. *The 2018 Annual Homeless Assessment Report (AHAR) to Congress: PART 2: Estimates of Homelessness in the United States*. Technical Report. The U.S. Department of Housing and Urban Development.
[31] Jennifer L Hill. 2011. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics* 20, 1 (2011), 217–240.
[32] Safwan Hossain, Andjela Mladenovic, and Nisarg Shah. 2020. Designing fairly fair classifiers via economic fairness notions. In *Proceedings of The Web Conference 2020*. 1559–1569.
[33] Marian Moser Jones. 2016. Does race matter in addressing homelessness? A review of the literature. *World Medical & Health Policy* 8, 2 (2016), 139–156.
[34] Jill Khadduri, Josh Leopold, Brian Sokol, and Brooke Spellman. 2010. Costs associated with first-time homelessness for families and individuals. *Available at SSRN 1581492* (2010).
[35] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent trade-offs in the fair determination of risk scores. In *Conference on Fairness, Accountability and Transparency*. PMLR.
[36] Amanda R Kube, Sanmay Das, and Patrick J Fowler. 2019. Fair and Efficient Allocation of Scarce Resources Based on Predicted Outcomes: Implications for Homeless Service Delivery. (2019).
[37] Zhuoshu Li, Kelsey Lieberman, William Macke, Sofia Carrillo, Chien-Ju Ho, Jason Wellen, and Sanmay Das. 2019. Incorporating compatible pairs in kidney exchange: A dynamic weighted matching model. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 349–367.
[38] David F Manlove and Colin TS Sng. 2006. Popular matchings in the capacitated house allocation problem. In *European Symposium on Algorithms*. Springer, 492–503.
[39] Matthew H Morton, Shannon Kugley, Richard Epstein, and Anne Farrell. 2020. Interventions for youth homelessness: A systematic review of effectiveness studies. *Children and Youth Services Review* 116 (2020), 105096.
[40] Heather Menzies Munthe-Kaas, Rigmor C Berg, and Nora Blaasvær. 2018. Effectiveness of interventions to reduce homelessness: A systematic review and meta-analysis. *Campbell Systematic Reviews* 14, 1 (2018), 1–281.
[41] Stephen Pfohl, Ben Marafino, Adrien Coulet, Fatima Rodriguez, Latha Palaniappan, and Nigam H Shah. 2019. Creating fair models of atherosclerotic cardiovascular disease risk. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 271–278.
[42] ProPublica. 2016. How We Analyzed the COMPAS Recidivism Algorithm. *ProPublica* (2016).
[43] Debra J Rog, C Scott Holupka, and Lisa C Patton. 2007. Characteristics and dynamics of homeless families with children. *Washington, DC: US Department of Health and Human Services* (2007).
[44] Alvin Roth. 2015. *Who gets what-and why: The hidden world of matchmaking and market design*. HarperCollins UK.
[45] Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. 2005. Pairwise kidney exchange. *Journal of Economic theory* 125, 2 (2005), 151–188.
[46] Hal R Varian. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9, 1 (1974), 63–91.
[47] Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. 2016. Using Social Networks to Aid Homeless Shelters: Dynamic Influence Maximization under Uncertainty.. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, Vol. 16. 740–748.