

# The Paradox of Prioritization in Public Sector Algorithms

ERINA SEH-YOUNG MOON, University of Toronto, Canada

SHION GUHA, University of Toronto, Canada

Public sector agencies perform the critical task of implementing the redistributive role of the State by acting as the leading provider of critical public services that many rely on. In recent years, public agencies have been increasingly adopting algorithmic prioritization tools to determine which individuals should be allocated scarce public resources. Prior work on these tools has largely focused on assessing and improving their fairness, accuracy, and validity. However, what remains understudied is how the structural design of prioritization itself shapes both the effectiveness of these tools and the experiences of those subject to them under realistic public sector conditions. In this study, we demonstrate the fallibility of adopting a prioritization approach in the public sector by showing how the underlying mechanisms of prioritization generate significant relative disparities between groups of intersectional identities as resources become increasingly scarce. We argue that despite prevailing arguments that prioritization of resources can lead to efficient allocation outcomes, prioritization can intensify perceptions of inequality for impacted individuals. We contend that efficiencies generated by algorithmic tools should not be conflated with the dominant rhetoric that efficiency necessarily entails “doing more with less” and we highlight the risks of overlooking resource constraints present in real-world implementation contexts.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**; • **Applied computing** → **Computing in government**.

Additional Key Words and Phrases: prioritization, algorithm, public sector, resource allocation, government

## ACM Reference Format:

Erina Seh-Young Moon and Shion Guha. 2026. The Paradox of Prioritization in Public Sector Algorithms. In *The 2026 ACM Conference on Fairness, Accountability, and Transparency (FAccT '26)*, June 25–28, 2026, Montreal, QC, Canada. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3805689.3812293>

## 1 Introduction

Public sector agencies perform the critical task of carrying out the redistributive role of the state and have increasingly adopted algorithmic prioritization tools to allocate essential resources effectively and efficiently [14, 21, 30, 31, 45, 53]. For example, in North America, homelessness service systems increasingly rely on predictive risk algorithms to assess the vulnerability of clients seeking housing supports in order to prioritize services to them, and public schools employ tools to determine which students should be prioritized to receive academic support to prevent school dropout [14, 30, 42, 53, 57]. Public sector prioritization algorithms are distinct from private sector decision-making algorithms as the tools are often applied to the most vulnerable members of society and states take an active role in controlling these populations [3, 14, 26].

Existing scholarship on public sector algorithms has largely focused on improving the models themselves, with an emphasis on enhancing model accuracy, fairness, and validity [25, 27, 46]. Yet, there have been growing calls within the algorithmic fairness community to move beyond this model-centered approach and the need

---

Authors' Contact Information: Erina Seh-Young Moon, University of Toronto, Toronto, Canada, [erina.moon@mail.utoronto.ca](mailto:erina.moon@mail.utoronto.ca); Shion Guha, University of Toronto, Toronto, Canada, [shion.guha@utoronto.ca](mailto:shion.guha@utoronto.ca).



This work is licensed under a Creative Commons Attribution 4.0 International License.

*FAccT '26, Montreal, QC, Canada*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2596-8/2026/06

<https://doi.org/10.1145/3805689.3812293>

to examine these tools as policy interventions in practice while considering their implementation context [27, 28, 52]. In this work, we respond to these calls by shifting the focus from the algorithmic model to the very structural mechanisms of prioritization that underpins public sector prioritization algorithms. We define structural mechanisms of prioritization as the design principles that specify how rank and weights structure the distribution of public goods. With this definition in mind, this paper provides a diagnostic analysis of the structural trade-offs that can necessarily emerge when adopting prioritization approaches. Specifically, our work asks the research question (RQ): **How does the structural design of public-sector prioritization algorithms shape resource allocation disparities?**

To address our research question, we examine two prioritization approaches that underpin many public sector resource allocation systems, a hierarchical prioritization approach and weighted prioritization approach, which we depict in Figure 1. These two approaches are commonly applied to guide the prioritization and allocation of public resources, including in healthcare, housing in homelessness systems, and educational services [14, 21, 30, 55]. We focus on the most basic version of each of these prioritization approaches as the aim of this work is to distill how the structural mechanisms underpinning these two methods can systematically shape allocation outcomes. We make the following contributions to the paper:

- We frame public sector prioritization algorithms as policy interventions [28] and examine the structural design of prioritization that underpins these models. We demonstrate how resource allocation disparities arise from inherent properties of these structures in Sections 4 and 5.
- We formalize how resource scarcity is a driver of inequality, particularly under a hierarchical prioritization regime in Section 5.2. We demonstrate that under this regime, disparities in the likelihood of receiving resources (log-ratio differences in receiving resources) can become arbitrarily large as resources become scarce.
- We demonstrate how different evaluation metrics introduced in Section 5.1 (absolute and log-ratio differences in resource allocation rates) for measuring the disparate impact of prioritization can reveal distinct experiential aspects of prioritization in Sections 5.2 and 5.3. We present a naïve example of prioritization in homelessness to highlight these distinctions in Section 5.4.

In the following section, we detail the related works that have motivated this work.

## 2 Related Works

In recent years, often driven by neoliberal and market-oriented values, public agencies have increasingly sought to categorize and classify individuals across multiple dimensions to prioritize resources toward targeted groups, particularly in contexts of resource scarcity to promote allocation efficiency [12, 14, 21, 41]. Public resources are not provided universally but instead, allocated through prioritization based on assessments of an individual's or entity's 'deservingness' or their 'risk level' [21]. Following the growing technical affordances of data-driven algorithmic systems, prioritization algorithms are being increasingly used to operationalize these assessments. Applications of these prioritization models in the public domain are far ranging: from criminal justice algorithms that predict areas at higher risk of criminal activity to guide the deployment of policing resources and homelessness algorithms that predict a person's risk of experiencing chronic homelessness to prioritize the allocation of housing-related services [31, 53, 59, 60].

The algorithmic fairness community has extensively studied these systems, where a large body of literature has focused on model-centric evaluations such as improving model accuracy, validity, reliability, and fairness [7, 9, 10, 15, 20, 24, 29]. However, in response to studies that highlight the unintended harms associated with the deployment of these tools [4, 14, 44, 47], Liu et al. [28] argue for the need to move away from adopting a model-focused paradigm and instead, towards framing algorithmic systems as policy interventions that must be evaluated in relation to their implementation context. We respond to these calls and shift our attention to the

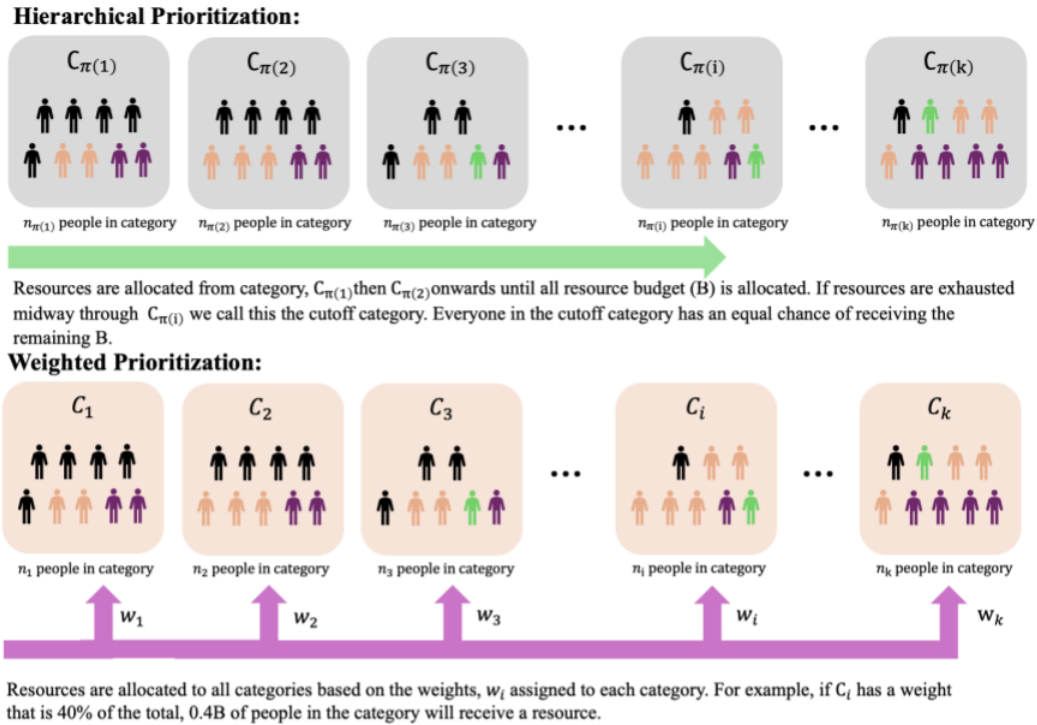


Fig. 1. Two Types of Prioritization Approaches for Resource Allocation: A strict hierarchical prioritization approach (1st row) and weighted prioritization approach (2nd row). Categories are composed of heterogeneous groups of individuals of diverse attributes

very structural mechanisms that underpin public sector prioritization algorithms. We examine how structural prioritization mechanisms that govern the ranking and allocation of public goods interact with varying conditions of resource scarcity to produce disparate outcomes for impacted groups. We also demonstrate how different measures of disparity can reveal different experiential effects of prioritization in the public sector [18, 28, 51].

### 3 Two Prioritization Approaches in the Public Sector

In this paper, we examine the structural mechanisms of two prioritization approaches: hierarchical and weighted prioritization, depicted in Figure 1. Although both approaches partition individuals into groups and allocate resources based on a person’s group membership (i.e., members of certain groups are prioritized for resources), they represent two distinct and commonly observed approaches to distributing public goods [54]. A hierarchical prioritization approach is a rank-based distribution approach that prioritizes resources to those with the highest need or deservingness [54]. Emergency care triage systems, child welfare systems, and public school systems often adopt this approach to identify critically ill patients, children at risk of experiencing maltreatment, or students at risk of dropping out to direct scarce resources to them [30, 48, 58]. In hierarchical prioritization, individuals are often ranked based on individual-level measurements, and lower-ranked categories may be excluded altogether from receiving resources [14]. A weighted prioritization approach, on the other hand, is a group-based distribution strategy that often divides people based on immutable characteristics, such as ethnicity, race, or gender (often historical markers of privilege and disadvantage), and is employed when representation

across groups is important [54]. For example, this approach is applied in public housing and vaccine rollouts to ensure no groups are completely excluded from receiving resources [37, 43]. While these two prioritization approaches are not exhaustive, we introduce them to provide a diverse illustration of how different prioritization structures give rise to resource allocation disparities.

To formalize the mechanics of prioritization, we introduce the following set up. We consider a population of  $N$  individuals in need of public resources, partitioned into  $K$  disjoint categories,  $C = \{C_1, C_2, \dots, C_K\}$ . These categories are determined by policymakers based on criteria such as risk scores, assessed level of need of a person, test scores, or other similar measures. Each category,  $C_i$  contains  $n_i$  individuals, where  $\sum_{i=1}^K n_i = N$ . We assume there is a fixed resource budget of  $B$  resources, where a person can only receive one unit at most. Oftentimes in the public sector there is resource scarcity where  $B < N$ .

### 3.1 Hierarchical Prioritization

The first row of Figure 1 depicts hierarchical prioritization. Under this scheme, categories are ordered according to a ranking criterion established by policymakers or an algorithm [21] and resources are allocated sequentially until the budget is exhausted. In cases where the remaining budget is insufficient to serve an entire category (we call this the "cutoff category"), we make the explicit assumption that everyone within that cutoff category is equally likely to receive the resource. We adopt this specific assumption as this allows us to isolate the impact of the hierarchical ranking structure on allocation outcomes (see Appendix A where we further discuss this assumption). While policymakers may assign identical ranks to multiple categories, we treat such instances as a single collapsed category for the purposes of this stylized model. When  $n_i = 1$  for all  $K$  categories, this is the special case where we have individual-level prioritization.

**DEFINITION 3.1 (HIERARCHICAL PRIORITIZATION).** *A hierarchical prioritization rule,  $\pi$ , ranks each category such that  $C_{\pi(1)} \succ C_{\pi(2)} \succ \dots \succ C_{\pi(K)}$ . Given a resource budget  $B$ , resources are allocated to the categories in order of rank from highest to lowest. The probability  $P_{\pi(i)}^H$  of an individual in category  $\pi(i)$  receiving the resource is defined as:*

$$P_{\pi(i)}^H = \begin{cases} 1, & \text{if } \sum_{j=1}^i n_{\pi(j)} \leq B, \\ \frac{B - \sum_{j=1}^{i-1} n_{\pi(j)}}{n_{\pi(i)}}, & \text{if } \sum_{j=1}^{i-1} n_{\pi(j)} < B < \sum_{j=1}^i n_{\pi(j)}, \\ 0, & \text{if } \sum_{j=1}^{i-1} n_{\pi(j)} \geq B. \end{cases} \quad (1)$$

Equation 1 formalizes the sequential assignment of resources shown in the first row of Figure 1. In the Figure, the cutoff category for hierarchical prioritization is the category where the green arrow ends (i.e.,  $C_{\pi(i)}$ ). In this prioritization regime, individuals in categories ranked higher than the cutoff category (categories to the left of  $C_{\pi(i)}$ ) are fully served and receive resources with a probability of 1. When the budget,  $B$  is partially exhausted in the cutoff category, individuals in this cutoff group receive resources with a probability equal to the remaining budget divided by the category size ( $n_{\pi(i)}$ ). Any individuals in lower ranked categories (those in categories to the right of the cutoff category in the Figure) have a 0 probability of receiving resources.

### 3.2 Weighted Prioritization

The second row of Figure 1 depicts weighted prioritization. This approach allocates resources across categories proportionally based on pre-assigned weights for each category. For example, if a category is assigned a weight that is 40% of the total sum of weights,  $0.4B$  resources will be assigned to the category (which means that  $0.4B$  people in the category will receive resources). Although in resource scarcity, not every individual in a category will necessarily receive a resource, this approach ensures that multiple categories receive an allocation

simultaneously, avoiding the "cutoff" effect seen in hierarchical prioritization where the budget may be exhausted before reaching lower-ranked groups. In the public sector, these weights,  $w_i$ , could represent the degree of need or deserving-ness for certain population groups (e.g., a fixed proportion of available homelessness shelters must go to Indigenous groups).

**DEFINITION 3.2 (WEIGHTED PRIORITIZATION).** *A weighted prioritization rule assigns a positive weight,  $w_i$  to each category,  $C_i$ . Given a resource budget  $B$ , the total amount of resources,  $R_i$  allocated to category  $i$  is*

$$R_i = \frac{w_i}{\sum_{j=1}^K w_j} \cdot B \quad (2)$$

The probability  $P_i^W$  of an individual in category  $i$  receiving the resource is defined as:

$$P_i^W = \min \left( 1, \frac{w_i \cdot B}{n_i \cdot \sum_{j=1}^K w_j} \right) \quad (3)$$

Note that as  $B$  becomes large, the resource allocated to a specific category,  $R_i$ , can become greater than the size of the category,  $n_i$ . For this stylized example, we apply a "no reallocation after saturation" rule: we cap the probability at 1, and any excess budget assigned to that category is considered unspent. This assumption allows us to isolate the effects of weighted prioritization in its most basic form. However, we acknowledge that in many real-world systems, excess resources are redirected to other categories, with reallocation mechanisms varying widely (E.g., weighted redistribution, uniform redistribution).

As illustrated by the purple arrows in the second row of Figure 1, every category receives a "slice" of the budget based on their assigned weight ( $w_i$ ). Equation 3 translates this category allocation weight into an individual probability. For example, if a category,  $i$ , of 10 people ( $n_i = 10$ ) is allocated 8 units ( $R_i = 8$ ), an individual has a  $P_i^W = 0.8$  probability of receiving resources. However, if this category instead has only 4 people ( $n_i = 4$ ), the "no reallocation after saturation" rule caps the probability at  $P_i^W = 1$ .

## 4 Structural Comparisons Between the Two Approaches

The structural design of the two prioritization approaches results in differences in how intra- and inter- group tensions arise within and between different categories, respectively.

### 4.1 Intra-category Tensions

Intra-category tensions arise when an individual's probability of receiving resources within a group decreases as the number of people in a category increases.

In the **hierarchical approach**, as seen in Eq. (1), for categories where the cumulative population is less than  $B$ , the probability of receiving resources for an individual,  $P_{\pi(i)}^H = 1$ . However, for a cutoff category,  $m$  where the remaining budget is insufficient to cover the entire group, the individual probability of a person receiving resources decreases as the group size grows. Formally, this decrease occurs at a decreasing rate:

$$\frac{\partial P_{\pi(i)}^H}{\partial n_{\pi(i)}} = -\frac{B - \sum_{j=1}^{i-1} n_{\pi(j)}}{(n_{\pi(i)})^2} < 0, \quad \frac{\partial^2 P_{\pi(i)}^H}{\partial n_{\pi(i)}^2} = \frac{2 \left( B - \sum_{j=1}^{i-1} n_{\pi(j)} \right)}{(n_{\pi(i)})^3} > 0 \quad (4)$$

Following Equations 4, we observe that for the cutoff category, adding a fixed number of people to the category reduces each person's probability of receiving resources more when that category is small than when it is large.

In the **weighted approach**, intra-category tensions are not confined within a single cutoff category as in the hierarchical approach but can be experienced across all categories. An increase in a category's group size results

in a decrease in the individual probability of a person receiving resources at a decreasing rate for unsaturated categories:

$$\frac{\partial P_i^W}{\partial n_i} = -\frac{w_i \cdot B}{n_i^2 \cdot \sum w_j} < 0, \quad \text{and} \quad \frac{\partial^2 P_i^W}{\partial n_i^2} = \frac{2w_i B}{n_i^3 \sum w_j} > 0 \quad (5)$$

Here, the impact of a person's chances of receiving resources following a change in their category's group size is scaled by the weight assigned to their category. While high  $w_i$  can favor the category in terms of total resources allocated, the marginal decrease in an individual's probability of receiving resources following an increase in category size is greater when the category is small than when it is large.

## 4.2 Inter-category Tensions

Inter-category tensions arise when changes in one category's group size affect the likelihood that individuals in another category receive resources.

In the **hierarchical approach**, for any category  $C_{\pi(k)}$  ranked lower than  $C_{\pi(i)}$ , an increase in the population of higher-ranked groups reduces the remaining resources for lower-ranked groups. In the marginal cutoff category, inter-category harm arises since:

$$\frac{\partial P_{\pi(k)}^H}{\partial n_{\pi(i)}} = -\frac{1}{n_{\pi(k)}} < 0 \quad \text{when } i < k \quad (6)$$

The derivative in Equation 6 shows that adding an individual to a higher-ranked category marginally reduces each member's probability of receiving resources in the cutoff category by  $\frac{1}{n_{\pi(k)}}$ . On the other hand, higher-ranked groups are unaffected by the population sizes of lower-ranked groups. Hierarchical prioritization approach thus creates a structural dynamic where low-priority groups are vulnerable to "budget exhaustion" caused by higher-priority groups.

Under our definition of weighted prioritization, which assumes 1) weights assigned to each category are fixed constant and 2) that the "no reallocation after saturation" rule holds (Definition 3.2), there is no inter-category tension in the **weighted approach**, because changes in the size of one category do not affect the probability of resource receipt for other categories. However, we note that when either of these assumptions is relaxed, inter-category tensions may arise as leftover resources can be redistributed to other categories. How these tensions emerge will depend on the excess resource redistribution policy.

## 5 Disparate Impact on Intersectional identities

In the previous section, we examined tensions that can arise within and between categories following hierarchical and weighted prioritization approaches. These tensions are simple, inherent structural properties of their respective allocation mechanisms. We now extend our analysis to examine how the design of hierarchical and weighted regimes can produce disparate impact for individuals with intersectional identities because classification of people into categories does not imply homogeneity within groups; instead, each category is often comprised of a heterogeneous population with diverse attributes. For example, in healthcare and child welfare systems, clients seeking services may be grouped based on risk level; however, within these groups, individuals may differ significantly by race, gender, or socioeconomic status [1, 2, 6, 14]. Prior works has shown how such categorized groups can systematically marginalize specific populations, particularly when these categories are strongly correlated with protected attributes such as poverty or race [11, 14, 16, 34].

### 5.1 Metrics to measure disparate impact

To formalize disparate impacts of prioritization on subgroups within categories, we consider a set of population subgroups:

$$\mathcal{S} = \{s_1, s_2, \dots, s_m\}$$

distributed across the  $K$  categories. In Figure 1, these subgroups are depicted by different colored figures depicted within categories. Let  $n_{i,s}$  denote the number of people in  $C_i$  who belong to subgroup,  $s$ . The total number of people in category  $n_i = \sum_{s \in \mathcal{S}} n_{i,s}$  and the total population of a specific subgroup is  $N_s = \sum_{i=1}^K n_{i,s}$ . Within any category  $C_i$ , we assume all individuals have an equal probability  $P_i$  of receiving a resource regardless of their subgroup membership. We intentionally assume no within-category discrimination to isolate structural effects and avoid introducing within-category bias.

To understand the disparate impacts of prioritization, we compare how the proportion of the total subgroup population expected to receive a resource differs between subgroups and across varying budgetary conditions. Below, we define the metrics we use to quantify subgroup allocation disparities.

**DEFINITION 5.1 (SUBGROUP RESOURCE RECEIPT RATE).** *The expected resource receipt rate for subgroup  $s$ , denoted  $G_s(B)$ , is the proportion of the total subgroup population that is expected to receive a resource:*

$$G_s(B) = \frac{\sum_{i=1}^K n_{i,s} \cdot P_i(B)}{N_s} \quad (7)$$

where  $P_i(B)$  represents the probability of an individual in category  $C_i$  receiving a resource given a total budget  $B$ .

**The Subgroup Resource Receipt Rate** represents the expected proportion of people in a particular subgroup who will receive resources given a resource budget of  $B$ . For example, if half of the people in subgroup  $s_1$  are in category 1 with a 100% chance of receiving resources, and the other half are in category 2 with a 20% chance of receiving resources, then 0.6 of those in subgroup,  $s_1$ , are expected to receive a resource.

To understand disparities in resource allocation between subgroups, we compare Subgroup Resource Receipt Rates across two metrics: log-ratio differences,  $\ln(RD(B))$ , and absolute difference,  $AD(B)$ , which we define below.

**DEFINITION 5.2 (ABSOLUTE DIFFERENCE IN SUBGROUP RESOURCE RECEIPT RATE).** *The absolute difference between two expected subgroup resource receipt rates (e.g., between subgroups  $s_1$  and  $s_2$ ) at budget  $B$  is defined as:*

$$AD(B) = |G_{s_1}(B) - G_{s_2}(B)| \quad (8)$$

**The Absolute Difference (AD)** captures how unequally resources are allocated between two subgroups under a given prioritization regime. For example, if 60% of individuals in subgroup,  $s_1$  and 40% of people in subgroup,  $s_2$  are expected to received resources then  $AD(B) = 0.2$ .

**DEFINITION 5.3 (LOG-RATIO DIFFERENCE IN SUBGROUP RESOURCE RECEIPT RATE).** *The ratio difference between two expected subgroup resource receipt rates at budget  $B$  is:*

$$RD(B) = \frac{G_{s_1}(B) + \epsilon}{G_{s_2}(B) + \epsilon} \quad (9)$$

where  $0 < \epsilon \ll 1$  is a small smoothing constant to ensure the ratio is well defined because  $RD(B)$  becomes undefined when denominator subgroup resource receipt rates equals 0. Taking the log of  $RD(B)$ , we define the log-ratio difference in subgroup resource receipt rates as:

$$\ln(RD(B)) = \ln\left(\frac{G_{s_1}(B) + \epsilon}{G_{s_2}(B) + \epsilon}\right) \quad (10)$$

*This measure provides a useful metric to understand differences in magnitude between the two groups' receipt rates. We note the absolute value of this log-ratio also provides a useful symmetric measure of relative disparity between subgroups. E.g., Taking the absolute value of  $\ln(RD(B))$  allows us to treat the ratio difference of  $RD(B) = 2$  the same as a ratio difference of  $RD(B) = \frac{1}{2}$ . Additionally, we note that the absolute value of the derivative of this Log-Ratio further provides a useful metric to capture the sensitivity of the prioritization scheme to budgetary fluctuations as this metric is also symmetric (explained further in Section 5.2 and 5.3).*

We apply a small epsilon in Definition 5.3 to ensure these quantities are well defined when the denominator subgroup resource receipt rate equals zero. In these instances, the resulting values should not be interpreted as an empirically meaningful measure of disparity. This differs from cases where the subgroup resource receipt rate is very low and near zero; in this instance disparity values reflect empirically substantive differences in expected resource allocation, which we discuss further below.

The **Log-Ratio Difference in Subgroup Receipt Rate** captures the proportional disparity of receiving resources between two subgroups. While the Absolute Difference measures the size of the gap between two Subgroup Receipt Rates, the Log-Ratio Difference focuses on how many times more likely one subgroup is to receive resources over another. For example, if subgroup  $s_1$  has an expected Resource Receipt Rate of 0.8 and subgroup  $s_2$  has a rate of 0.4, then  $RD(B) = 2$ , meaning those in subgroup  $s_1$  are twice as likely to receive resources as those in subgroup  $s_2$ . The logarithmic transformation (Log-Ratio Difference in Subgroup Resource Receipt Rate) provides mathematical symmetry, ensuring that an advantage for one subgroup is given the same weight as an equivalent advantage for the other. For instance, if  $s_1$  is twice as likely as  $s_2$  to receive resources, the Log-Ratio is approximately 0.69; if the situation were reversed and  $s_2$  were twice as likely, the value would be  $-0.69$ . Taking the absolute value of this metric allows us to measure the intensity of the disparity regardless of which subgroup is more advantaged.

## 5.2 Hierarchical Prioritization can Produce Explosive Subgroup Disparities in High Resource Scarcity

Beyond the inter- and intra-category tensions discussed in Section 4, hierarchical prioritization can give rise to intersectional disparities in resource receipt when categories differ in their composition of intersectional subgroups. When the resource budget,  $B$ , falls within the cutoff category,  $C_{\pi(m)}$  (i.e., the budget cannot serve all the individuals in the cutoff category), the marginal benefit to subgroup  $s$  as the budget,  $B$  changes is as follows (see Appendix B for the derivation):

$$\frac{dG_s^H}{dB} = \frac{1}{N_s} \cdot \frac{n_{\pi(m),s}}{n_{\pi(m)}} \quad (11)$$

Equation 11 demonstrates that when the budget increases, a subgroup's chances of receiving resources depend on the subgroup's density within the cutoff category and the subgroup's size across all categories. In short, if the subgroup is heavily concentrated in the cutoff category, a unit increase in the resource budget will substantially increase the subgroup's Resource Receipt Rate. However, if the subgroup population is spread thin across many categories, a unit increase in the budget will only yield a modest increase in the subgroup's Resource Receipt Rate.

We observe the most interesting findings when we compare the Absolute Difference in Resource Receipt Rates,  $AD^H(B)$  and Log-Ratio Differences in Resource Receipts rates,  $\ln(RD^H(B))$  as they respond to changes in  $B$ . In hierarchical prioritization,  $AD^H(B)$  moves in a piecewise linear fashion as we show below:

$$\frac{d}{dB} AD^H(B) = \text{sgn}(G_{s_1}^H - G_{s_2}^H) \left( \frac{n_{\pi(m),s_1}}{N_{s_1} \cdot n_{\pi(m)}} - \frac{n_{\pi(m),s_2}}{N_{s_2} \cdot n_{\pi(m)}} \right) \quad (12)$$

The derivative in Equation 12 is undefined when  $G_{s_1}^H = G_{s_2}^H$  and its sign flips as the budget crosses this point. As the budget increases to cover more categories,  $AD^H(B)$  changes at a constant rate within each category. The

rate is dependent on the difference in the subgroups' proportional representation within the cutoff category relative to each subgroup's total size. For example, if a subgroup  $s_1$  has a higher receipt rate than subgroup  $s_2$  (I.e.,  $G_{s_1}^H > G_{s_2}^H$ ) and there is a higher proportion of  $s_1$  in the cutoff category compared to  $s_2$ , then we would expect to see the gap between the Resource Receipt Rates for the two subgroups to increase as the budget increases. Conversely, if a subgroup  $s_2$  has a higher receipt rate than subgroup  $s_1$  (I.e.,  $G_{s_1}^H < G_{s_2}^H$ ) and there are a higher proportion of  $s_1$  in the cutoff category compared to  $s_2$ , then we would expect to see the gap between the Resource Receipt Rates for the two subgroups to decrease as the budget increases.

While Equation 12 shows that  $AD^H(B)$  moves in a piecewise linear fashion,  $RD^H(B)$  exhibits greater volatility and can create explosive disparities between subgroups in hierarchical approaches. The sensitivity of the  $RD^H(B)$  is best captured by observing the absolute value of the derivative of the log-ratio differences (see Definition 5.3):

$$\left| \frac{d}{dB} \ln(RD^H(B)) \right| = \left| \frac{G_{s_1}^H(B)}{G_{s_1}^H(B) + \epsilon} - \frac{G_{s_2}^H(B)}{G_{s_2}^H(B) + \epsilon} \right| = \left| \frac{\frac{n_{\pi(m),s_1}}{N_{s_1} \cdot n_{\pi(m)}}}{G_{s_1}^H(B) + \epsilon} - \frac{\frac{n_{\pi(m),s_2}}{N_{s_2} \cdot n_{\pi(m)}}}{G_{s_2}^H(B) + \epsilon} \right| \quad (13)$$

As noted in Definition 5.3, the absolute value of the derivative of  $\ln(RD^H(B))$  captures how sensitive the relative likelihood of receiving resources is for one subgroup compared to another as the resource budget changes. Equation 13 shows that the derivative is calculated by the inverse of the subgroup receipt rates. Our findings imply that when 1) the resource budget falls within the cutoff category, 2) subgroup resource receipt rates that are being compared are positive, and 3) when there is high resource scarcity such that a subgroup's resource receipt rate is very low and near zero, then even slight budgetary shifts can create explosive changes in  $RD^H(B)$ .

In short,  $AD^H(B)$  changes steadily as the budget grows but  $\ln(RD^H(B))$  can fluctuate sharply. This happens because when a subgroup receives very few resources, even a small budget increase can dramatically change their relative chances of getting resources. For example, consider the scenario where  $G_{s_1}^H(B) = 0.4$  and  $G_{s_2}^H(B) = 0.01$ . Here the Absolute Difference is  $AD^H(B) = 0.39$  while the Log-Ratio Difference is  $\ln(RD^H(B)) = 3.69$ . Now, if a small budget increase adds just 0.01 to each rate ( $G_{s_1}^H(B) = 0.41$  and  $G_{s_2}^H(B) = 0.02$ ), the Absolute Difference will remain at 0.39, but the Log-Ratio Difference drops sharply to  $\ln(RD^H(B)) = 3.02$ . The divergence between  $AD^H(B)$  and  $\ln(RD^H(B))$  as the budget changes is important because for impacted individuals who are hoping to receive resources,  $\ln(RD^H(B))$  may be a more pertinent measure that reflects their experiential relative likelihood of receiving resources.

### 5.3 Weighted Prioritization Avoids Explosive Subgroup Disparities

We now examine the weighted prioritization approach. By assigning a nonzero weight ( $w_i > 0$ ) to all categories, this approach provides a way to ensure that no category is entirely excluded from receiving resources. In weighted prioritization, as the resource budget grows, the resource share allocated to each category will increase. Eventually, the budget will become large enough to allocate resources to every individual within a specific category ( $n_i$ ), reaching what we define as the saturation threshold. Because each category may have a different size and weight, the saturation thresholds will occur at different budgetary thresholds, E.g.,  $\{B_{sat,1}, B_{sat,2} \dots B_{sat,k}\}$ . In the following paragraphs, we denote  $\mathcal{A}(B)$  as the set of categories that are not yet saturated (i.e.,  $R_i < n_i$ ) and  $\mathcal{F}(B)$  as the set of saturated categories (i.e.,  $R_i \geq n_i$ ), the group receipt rate for subgroup  $s$  is (note for our stylized model, we assume weights are static):

$$G_s^W(B) = \frac{1}{N_s} \left( \sum_{i \in \mathcal{F}(B)} n_{i,s} + \sum_{i \in \mathcal{A}(B)} B \left[ \frac{w_i \cdot n_{i,s}}{n_i \cdot \sum_{j=1}^K w_j} \right] \right) \quad (14)$$

Equation 14 shows how the overall Resource Receipt Rate for a subgroup is determined under Weighted Prioritization. We first count all subgroup members in the saturated categories ( $\mathcal{F}(B)$ ), where every individual in

these categories is served a resource. To this, we add the expected number of subgroup members who receive resources in the unsaturated categories ( $\mathcal{A}(B)$ ). Finally, we divide this total by the subgroup's population size ( $N_s$ ).

Similar to hierarchical prioritization (Equation 12), the absolute difference in resource receipt rates,  $AD^W(B)$  moves in a piecewise fashion as the budget changes:

$$\frac{d}{dB}AD^W(B) = \text{sgn}\left(G_{s_1}^W(B) - G_{s_2}^W(B)\right) \left[ \sum_{i \in \mathcal{A}(B)} \frac{w_i}{\sum_{j=1}^K w_j} \left( \frac{n_{i,s_1}}{n_i N_{s_1}} - \frac{n_{i,s_2}}{n_i N_{s_2}} \right) \right] \quad (15)$$

Equation 15 shows that the change in the Absolute Difference in Resource Receipt Rate as the budget increases depends on categories that are not yet saturated. Each unsaturated category can either increase or decrease the gap, depending on which subgroup is more strongly represented in that category relative to the subgroup's total population. The magnitude of this effect is scaled by the weight assigned to the category; a higher weight means that the category has a larger influence on whether  $AD^W(B)$  widens or narrows.

The greatest difference between the weighted and hierarchical approaches can be observed in the Log-Ratio Difference of Resource Receipt Rates,  $\ln(RD^W(B))$ . When subgroups being compared have positive resource receipt rates, we no longer observe the explosive increase in log-ratio differences we observed in hierarchical prioritization. In weighted prioritization, prior to any category becoming saturated, the Log-Ratio Difference in Subgroup Receipt Rates,  $\ln(RD^W(B))$  is:

$$\ln(RD^W(B)) = \ln\left(\frac{G_{s_1}(B) + \epsilon}{G_{s_2}(B) + \epsilon}\right) = \ln\left(\frac{\frac{B}{N_{s_1} \sum w_j} \sum_{i \in \mathcal{A}(B)} \frac{w_i n_{i,s_1}}{n_i} + \epsilon}{\frac{B}{N_{s_2} \sum w_j} \sum_{i \in \mathcal{A}(B)} \frac{w_i n_{i,s_2}}{n_i} + \epsilon}\right) \quad (16)$$

When the budget is larger relative to the smoothing term  $\epsilon$  then  $B$  approximately cancels out, making the Log-Ratio Difference nearly constant. This implies that, as long as no category becomes saturated, the relative likelihood of receiving resources for one subgroup compared to another does not fluctuate greatly as the budget increases. As the budget continues to increase, and once one or more categories become saturated, then  $RD(B) \rightarrow 1$  (equivalently  $\ln(RD(B)) \rightarrow 0$ ) (reflecting that the subgroups are moving towards parity). These findings hold under the "no reallocation after saturation" rule. If this rule is no longer upheld and excess resources can be allocated to other categories, then we could observe lower absolute and log-ratio disparities (however, this is contingent on how excess resources are redistributed).

#### 5.4 Visualizing Resource Allocation Disparities between Subgroups: An Example in Homelessness

In this section, we visualize the disparate impact of prioritization in the context of homelessness where there is high demand across North American cities for emergency shelter beds and affordable permanent housing. In response to rising demand for housing-related services, federal governments in North America have adopted prioritization approaches to allocate housing-related resources, including emergency shelter spaces [13, 35].

We pull publicly available shelter population data from a large North American city [38, 39] collected in October 2024 to obtain counts of families and single adults where this data can be disaggregated by their refugee claimant status. Ethnographic scholarship suggests that families are, in the most general sense, perceived as more vulnerable than single adults and are often prioritized for services [40]. As a result, in this example, we consider two regimes that prioritize families over single adults for services: (1) a hierarchical prioritization regime where families are prioritized over single adults for access to permanent housing. And (2) a weighted prioritization regime where we arbitrarily assign weights of 0.7 and 0.3 to Families ( $C_1$ ) and Single Adults ( $C_2$ ), respectively. The distribution of refugee and non-refugee subgroups within these categories is shown in the top-left panel of Figure 2.

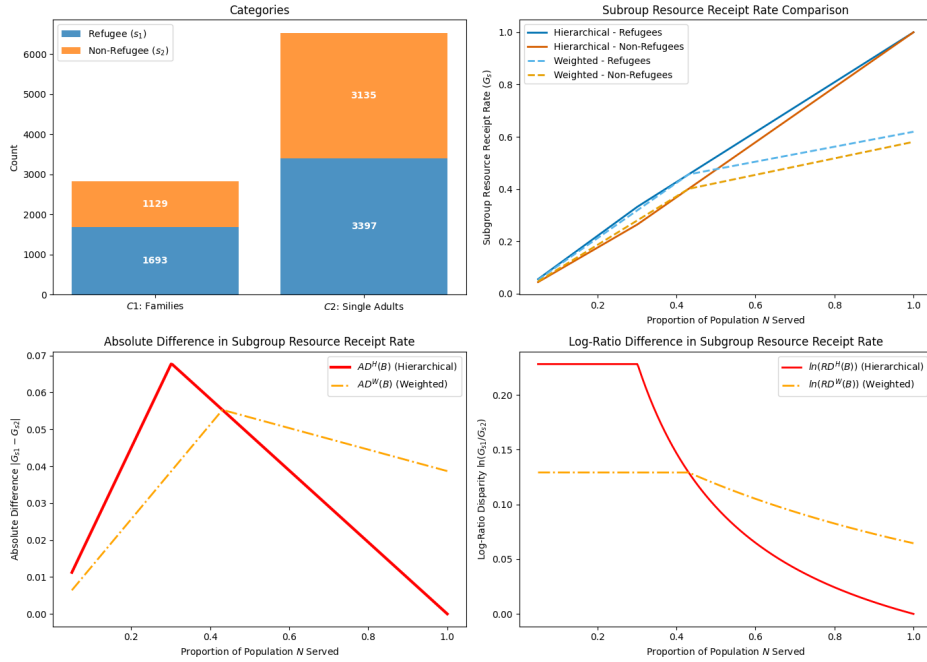


Fig. 2. Comparing Subgroup Resource Receipt Rates between Two Unhoused Groups, Families and Single Adults for Weighted and Hierarchical Prioritization: (Top-Left) Population distribution across categories. (Top-Right) Subgroup resource receipt rate ( $G_s$ ) comparisons as a function of total population served. (Bottom-Left) Absolute Difference ( $AD(B)$ ) in subgroup resource receipt rates between refugees and non-refugees. (Bottom-Right) Log-Ratio differences ( $\ln(RD(B))$ ) in subgroup resource receipt rates between refugees and non-refugees

As seen in the top-right panel of Figure 2, regardless of the prioritization approach, refugees have a higher Subgroup Resource Receipt Rate ( $G_s$ ) than non-refugees. Because resources are allocated sequentially under hierarchical prioritization (solid lines),  $G_s^H$  approaches 1 as  $B \rightarrow N$  (i.e., when the budget serves the entire  $N$  population). However, because we apply static weights and the "no reallocation after saturation rule" in the weighted prioritization approach (dashed lines), as the assignment of resources to a category reaches saturation (i.e.,  $R_i = n_i$ ),  $G_s^W$  levels off and does not reach 1 as  $B \rightarrow N$  since there are unused resources.

The bottom left panel illustrates how the absolute difference in receipt rates for the subgroups varies with changes in resource availability. Regardless of the prioritization approach, we observe that  $AD(B)$  moves in a piecewise linear manner and is bounded. Furthermore, as  $B$  increases from 0, the budget is being allocated to the subgroup that holds a higher allocation probability in this regime (in our example, refugees), increasing  $AD(B)$ . However, as the budget begins to be allocated to the subgroup with the lower allocation probability in our regime (in our example, non-refugees),  $AD(B)$  begins to decrease. Again, due to static weights and the "no reallocation after saturation rule," in weighted prioritization,  $AD^W(B)$  never reaches 0 at full resource availability.

In the bottom right panel, we see more volatility in  $\ln(RD^H(B))$  in hierarchical prioritization compared to  $AD(B)$ . In high resource scarcity, the magnitude of the receipt rate disparity is at its peak; and  $\ln(RD^H(B))$  remains flat when resources are allocated only to the first category. On the other hand,  $\ln(RD^W(B))$  remains flat for a longer budget interval as resources are distributed across multiple categories simultaneously.  $\ln(RD^W(B))$

only begins to curve downwards the moment the higher weighted category, Families ( $C_1$ ), are all assigned resources.

Our study’s most compelling findings emerge when we compare  $\ln(RD^H(B))$  and  $AD^H(B)$  under hierarchical prioritization. While  $AD^H(B)$  remains bounded and relatively low across all budget levels,  $\ln(RD^H(B))$  reveals severe disparities when there is high resource scarcity. From a policymaker’s perspective, the relatively stable, low values in  $AD^H(B)$  might suggest that subgroup disparities are acceptable. However, from the lived perspective of an unhoused person, the impact of prioritization could be experienced through  $\ln(RD^H(B))$ . In our example, at low budget levels, the relative likelihood of receiving housing is much lower for non-refugees. Consider when only 5% of the shelter population will receive housing (a realistic assumption for the City [39]). The resource receipt rate for a non-refugee,  $G_{Nonrefugee}(B = 468) = 0.044$  while for a refugee,  $G_{Refugee}(B = 468) = 0.055$ . In absolute nominal terms, the differences are small. However, from the perspective of a non-refugee seeking housing, refugees would appear to have 26% higher chance of obtaining housing. In contrast, by distributing resources across categories, the weighted prioritization can dampen this disparity to 13%. Overall, our example suggests that outcomes for all subgroups worsens with increasing housing scarcity. However, individuals in non-prioritized subgroups may feel a greater comparative disadvantage.

## 6 Discussion

In this work, we examined how the structural design of two common prioritization approaches applied in the public sector (weighted and hierarchical) can lead to resource allocation disparities between and within groups (Section 4). We further demonstrate how subgroup allocation disparities can arise from the inherent properties of the prioritization regime itself (Section 5). Our findings from Section 5.2 in particular highlight the implications of applying hierarchical prioritization to allocate public goods under conditions of extreme budget scarcity. In these contexts, we observe that resource allocation disparities become acutely sensitive to the proportional representation of subgroup population sizes. Although these effects can be mitigated under weighted prioritization (Section 5.3), overall, we observe relative allocation disparities between population groups arise from the structural mechanisms of prioritization. Our findings are important because we show distributive disparities emerge as direct consequences of intentional and normative policy choices on who to prioritize or not.

The distributive disparities identified in our work carry important social and policy implications for the allocation of scarce public resources using prioritization algorithms. Two key implications arise from our work: 1) we highlight the often overlooked trade-offs between efficiency and equity-related concerns in popular policy discourse on algorithmic tools, and 2) how allocation disparities may be experienced differently depending on whether they are viewed from a policymaker’s perspective or from the impacted individual. Predictive risk prioritization algorithms (a variant of hierarchical prioritization) have long been touted for their ability to target and deliver social services efficiently [12]. In recent years, amidst the growing technical affordances of AI tools, there is renewed interest in leveraging such computational models to deliver greater governmental efficiencies [17, 33, 36, 50]. Oftentimes, governmental "efficiency" is conflated with "doing more with less," as witnessed by the cuts made under the US Department of Government Efficiency (DOGE) in 2025 [19]. Our findings demonstrate that achieving these efficiency gains can invoke critical normative trade-offs. As a result, efficient prioritization of resources should not be necessarily conflated with the reduction of resource availability as this can come at the cost of stark allocation disparities when resources are scarce.

Stone [54] argues that distribution processes can strengthen or undermine social cohesion. Our work does not seek to de-legitimize prioritization but rather to clarify the tradeoffs prioritization structures create under scarcity. We hope our work encourages renewed scrutiny of the role that algorithmic prioritization tools should play in public sector resource allocation tasks and their societal impact. Prioritization algorithms primarily address the question of "who" should receive scarce public goods. In this work, we make explicit how different prioritization

structures give rise to disparate impacts, and how these impacts depend on the perspective from which they are evaluated. We introduce and compare two metrics, the absolute difference in subgroup receipt rates ( $AD(B)$ ) and the log-ratio difference in subgroup receipt rates ( $\ln(RD(B))$ ), to demonstrate how the prioritization regime is experienced/perceived depends on the lens through which it is measured. From a top-down, social planner's perspective, policymakers may deem a prioritization regime acceptable even under high resource scarcity due to small nominal differences in  $AD(B)$ . However, from the perspective of impacted individuals, ratio comparisons ( $\ln(RD(B))$ ) may be a more appropriate metric as individuals experience resource allocation as competition within their own subgroup and assess their relative likelihood of receiving resources in comparison to other subgroups.

The approaches we scrutinize contrast with other non-algorithmic allocation strategies implemented in the public sector, such as lotteries or first-come-first-served (FCFS) used in the provision of resources, for example, the provision of admission spots in charter schools [8]. These approaches can arguably offer a stronger form of procedural fairness around "how" scarce goods should be allocated [54]. And yet, when resources are scarce and when some individuals face acutely higher needs over others, adopting a non-prioritarian allocation approach may be deemed ethically or politically infeasible [56].

Returning to our example in Section 5.4, which visualizes the disparate impact of prioritization under varying resource availability, we observe that our findings succinctly capture the experiential disparities documented in ethnographic studies on how different population groups experience homelessness in North America. Previously, Passaro [40] argued that, shaped by normative societal biases regarding the worth and deservingness of different groups, North American single adult men have been placed at the bottom of every rung in receiving housing-related support. Passaro claimed that policies that "ignore men are helping to create increasingly desperate and violent men. This is not the result that anybody had in mind in attempting to protect women and children; it is the unintended outcome of what are well-meaning efforts to help." Although homelessness systems have undergone substantial changes since Passaro's writing, particularly following the adoption of centralized prioritization approaches (often referred to as coordinated entry and access systems [13]), the paradoxical effects of prioritization remain. Through our work, we demonstrate that adopting prioritization under high resource scarcity will inevitably lead to the creation of worse-off groups, and the magnitude of disparity between groups is directly tied to the structure of prioritization itself.

## 7 Limitations and Future Work

This work examines how disparities can arise from the structural design of two prioritization approaches in their most basic form by making several simplifying assumptions (e.g., assuming static weights and a "no reallocation after saturation" rule in weighted prioritization, and equal chances of resource receipt for all individuals within the cutoff category in hierarchical prioritization). Furthermore, we focused on one round of resource allocation through prioritization and did not consider how disparities can emerge under other distribution policies, such as first-come-first-serve or random lotteries. In practice, public-sector prioritization algorithms operate in different policy and organizational contexts, where multiple distributive policies and algorithms operate simultaneously [47, 55]. These factors may amplify or diminish the disparities we identify in our paper. For example, Saxena and Guha [49] show that once an allegation of child abuse is made and a family engages with the child welfare system, multiple prioritization algorithms may be deployed at multiple stages of a family's journey in the system. Thus, the extent to which the disparities we observe in our work manifest in practice will depend on specific domain implementation contexts and the discretionary decision-making of street-level bureaucrats [5, 22, 23, 32]. Lastly, while this work focused on the structural mechanics of "who" gets resources in prioritization algorithms, future work can examine how disparities also emerge when algorithms determine "what" type of resources are allocated and "how" they are delivered to individuals.

## 8 Conclusion

In this study, we examined how the structural design of prioritization (weighted and hierarchical prioritization) commonly applied in public sector prioritization algorithms give rise to disparate allocation outcomes for different population groups of intersectional identities. We introduced two evaluation metrics to assess the disparate impact of prioritization (log-ratio differences and absolute differences) and highlight how the metrics reveal different perspectives on how a prioritization policy is experienced and perceived.

## 9 Generative AI Usage Statement

Generative AI (Gemini and Chatgpt) was used to check the text for typographical errors and assist in the formatting of Figure 2. The authors reviewed all these suggestions generated by the AI models.

## Acknowledgments

This research was supported by the NSERC Discovery Early Career Researcher Grant RGPIN-2022-04570 and University of Toronto School of Cities Catalyst Grant. Opinions, findings, and conclusions expressed in this paper are those of the authors. We sincerely thank our anonymous reviewers whose suggestions and comments helped improve this manuscript.

## References

- [1] Julia Angwin Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. *Machine Bias*. Retrieved August 30, 2022 from <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [2] Ruha Benjamin. 2019. *Race after technology: abolitionist tools for the New Jim Code*. Polity Press, Cambridge, UK.
- [3] Khiara M. Bridges. 2017. *The Poverty of Privacy Rights*. Stanford University Press, Stanford, CA, USA.
- [4] Adriane Chapman, Philip Grylls, Pamela Ugwu-dike, David Gammack, and Jacqui Ayling. 2022. A Data-driven analysis of the interplay between Criminological theory and predictive policing algorithms. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAcCT '22). Association for Computing Machinery, New York, NY, USA, 36–45. doi:10.1145/3531146.3533071
- [5] Hao-Fei Cheng, Logan Stapleton, Anna Kawakami, Venkatesh Sivaraman, Yanghui Cheng, Diana Qing, Adam Perer, Kenneth Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. How Child Welfare Workers Reduce Racial Disparities in Algorithmic Decisions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 162, 22 pages. doi:10.1145/3491102.3501831
- [6] Lingwei Cheng, Cameron Drayton, Alexandra Chouldechova, and Rhema Vaithianathan. 2024. Algorithm-Assisted Decision Making and Racial Disparities in Housing: A Study of the Allegheny Housing Assessment Tool. arXiv:2407.21209 [cs.HC] <https://arxiv.org/abs/2407.21209>
- [7] Alexandra Chouldechova. 2017. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data* 5, 2 (2017), 153–163. doi:10.1089/big.2016.0047
- [8] Sarah Cohodes and Susha Roy. 2025. Thirty Years of Charter Schools: What Does Lottery-Based Research Tell Us? *Journal of School Choice* 19, 1 (2025), 8–49. doi:10.1080/15582159.2024.2379644
- [9] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. 2017. Algorithmic Decision Making and the Cost of Fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*. Association for Computing Machinery, New York, NY, USA, 797–806. doi:10.1145/3097983.3098095
- [10] Amanda Coston, Anna Kawakami, Haiyi Zhu, Ken Holstein, and Hoda Heidari. 2023. A Validity Perspective on Evaluating the Justified Use of Data-driven Decision-making Algorithms. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*. IEEE Computer Society, Los Alamitos, CA, USA, 690–704. doi:10.1109/SaTML54575.2023.00050
- [11] Courtney Cronley. 2022. Invisible intersectionality in measuring vulnerability among individuals experiencing homelessness – critically appraising the VI-SPDAT. *Journal of Social Distress and Homelessness* 31, 1 (2022), 23–33. doi:10.1080/10530789.2020.1852502
- [12] Lina Dencik, Joanna Redden, Arne Hintz, and Harry Warne. 2019. The ‘golden view’: data-driven governance in the scoring society. *Internet Policy Review* 8, 2 (2019). <https://policyreview.info/articles/analysis/golden-view-data-driven-governance-scoring-society>
- [13] John Ecker, Molly Brown, Tim Aubry, Katherine Francombe Pridham, and Stephen W. Hwang. 2022. Coordinated Access and Coordinated Entry System Processes in the Housing and Homelessness Sector: A Critical Commentary on Current Practices. *Housing Policy Debate* 32, 6 (Nov 2022), 876–895. doi:10.1080/10511482.2022.2058580
- [14] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin’s Press.

- [15] Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. 2021. The (Im)possibility of fairness: different value systems require different mechanisms for fair decision making. *Commun. ACM* 64, 4 (March 2021), 136–143. doi:10.1145/3433949
- [16] Marissa Gerchick, Tobi Jegede, Tarak Shah, Ana Gutierrez, Sophie Beiers, Noam Shemtov, Kath Xu, Anjana Samant, and Aaron Horowitz. 2023. The Devil is in the Details: Interrogating Values Embedded in the Allegheny Family Screening Tool. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 1292–1310. doi:10.1145/3593013.3594081
- [17] Gov.uk. 2025. *Ground-breaking use of AI saves taxpayers' money and delivers greater government efficiency*. Retrieved Jan 11, 2026 from <https://www.gov.uk/government/news/ground-breaking-use-of-ai-saves-taxpayers-money-and-delivers-greater-government-efficiency>
- [18] Ben Green. 2020. The false promise of risk assessments: epistemic reform and the limits of fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT\* '20). Association for Computing Machinery, New York, NY, USA, 594–606. doi:10.1145/3351095.3372869
- [19] The White House. 2025. *ESTABLISHING AND IMPLEMENTING THE PRESIDENT'S "DEPARTMENT OF GOVERNMENT EFFICIENCY"*. Retrieved Sept 2, 2025 from <https://www.whitehouse.gov/presidential-actions/2025/01/establishing-and-implementing-the-presidents-department-of-government-efficiency/>
- [20] Abigail Z. Jacobs, Su Lin Blodgett, Solon Barocas, Hal Daumé, and Hanna Wallach. 2020. The Meaning and Measurement of Bias: Lessons from Natural Language Processing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT\* '20). Association for Computing Machinery, New York, NY, USA, 706. doi:10.1145/3351095.3375671
- [21] Rebecca Ann Johnson and Simone Zhang. 2022. What is the Bureaucratic Counterfactual? Categorical versus Algorithmic Prioritization in U.S. Social Policy. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 1671–1682. doi:10.1145/3531146.3533223
- [22] Naveena Karusala, Sohini Upadhyay, Rajesh Veeraraghavan, and Krzysztof Z. Gajos. 2024. Understanding Contestability on the Margins: Implications for the Design of Algorithmic Decision-making in Public Services. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 478, 16 pages. doi:10.1145/3613904.3641898
- [23] Anna Kawakami, Venkatesh Sivaraman, Hao-Fei Cheng, Logan Stapleton, Yanghui Cheng, Diana Qing, Adam Perer, Zhiwei Steven Wu, Haiyi Zhu, and Kenneth Holstein. 2022. Improving Human-AI Partnerships in Child Welfare: Understanding Worker Practices, Challenges, and Desires for Algorithmic Decision Support. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (CHI '22). Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3491102.3517439
- [24] Jon Kleinberg. 2018. Inherent Trade-Offs in Algorithmic Fairness. *SIGMETRICS Perform. Eval. Rev.* 46, 1 (June 2018), 40. doi:10.1145/3292040.3219634
- [25] Ramaravind Kommiya Mothilal, Shion Guha, and Syed Ishtiaque Ahmed. 2024. Towards a Non-Ideal Methodological Framework for Responsible ML. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 477, 17 pages. doi:10.1145/3613904.3642501
- [26] Karen Levy, Kyla E. Chasalow, and Sarah Riley. 2021. Algorithms and Decision-Making in the Public Sector. *Annual Review of Law and Social Science* 17, 1 (2021), 309–334. doi:10.1146/annurev-lawsocsci-041221-023808
- [27] Lydia T. Liu, Solon Barocas, Jon Kleinberg, and Karen Levy. 2024. On the actionability of outcome prediction. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence (AAAI'24/IAAI'24/EAAI'24, Vol. 38)*. AAAI Press, 22240–22249. doi:10.1609/aaai.v38i20.30229
- [28] Lydia T. Liu, Inioluwa Deborah Raji, Angela Zhou, Luke Guerdan, Jessica Hullman, Daniel Malinsky, Bryan Wilder, Simone Zhang, Hammaad Adam, Amanda Coston, Ben Laufer, Ezinne Nwankwo, Michael Zanger-Tishler, Eli Ben-Michael, Solon Barocas, Avi Feller, Marissa Gerchick, Talia Gillis, Shion Guha, Daniel Ho, Lily Hu, Kosuke Imai, Sayash Kapoor, Joshua Loftus, Razieh Nabi, Arvind Narayanan, Ben Recht, Juan Carlos Perdomo, Matthew Salganik, Mark Sendak, Alexander Tolbert, Berk Ustun, Suresh Venkatasubramanian, Angelina Wang, and Ashia Wilson. 2025. Bridging Prediction and Intervention Problems in Social Systems. arXiv:2507.05216 [cs.LG] <https://arxiv.org/abs/2507.05216>
- [29] 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)*. Association for Computing Machinery, New York, NY, USA, 1095–1105. doi:10.1145/3531146.3533171
- [30] Kelly McConvey, Shion Guha, and Anastasia Kuzminykh. 2023. A Human-Centered Review of Algorithms in Decision-Making in Higher Education. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 223, 15 pages. doi:10.1145/3544548.3580658
- [31] Erina Seh-Young Moon and Shion Guha. 2024. A Human-Centered Review of Algorithms in Homelessness Research. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 870, 15 pages. doi:10.1145/3613904.3642392

- [32] Erina Seh-Young Moon, Devansh Saxena, Dipto Das, and Shion Guha. 2025. The Datafication of Care in Public Homelessness Services. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 829, 16 pages. doi:10.1145/3706598.3713232
- [33] Erina Seh-Young Moon, Matthew Tamura, Angelina Zhai, Nuzaira Habib, Behnaz Shirazi, Altaf Kassam, Devansh Saxena, and Shion Guha. 2026. The Promises and Perils of using LLMs for Effective Public Services. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 696, 21 pages. doi:10.1145/3772318.3790297
- [34] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453. arXiv:https://www.science.org/doi/pdf/10.1126/science.aax2342 doi:10.1126/science.aax2342
- [35] Government of Canada. 2023. *About Reaching Home: Canada's Homelessness Strategy*. Retrieved August 31, 2023 from https://www.infrastructure.gc.ca/homelessness-sans-abri/index-eng.html
- [36] Government of Canada. 2025. *Canada partners with Cohere to accelerate world-leading artificial intelligence*. Retrieved Sept 2, 2025 from https://www.canada.ca/en/innovation-science-economic-development/news/2025/08/canada-partners-with-cohere-to-accelerate-world-leading-artificial-intelligence.html
- [37] City of Toronto. 2021. *Homelessness Solutions Service Plan*. Retrieved aug 17, 2024 from https://www.toronto.ca/wp-content/uploads/2021/11/9885-Attachment-2-Homelessness-Solutions-Service-Plan.pdf
- [38] City of Toronto. 2025. Street Needs Assessment. Retrieved jan 4, 2026 from https://www.toronto.ca/community-people/housing-shelter/homeless-help/about-torontos-shelter-system/street-needs-assessment/
- [39] City of Toronto. 2026. Shelter System Flow Data. Retrieved jan 1, 2026 from https://www.toronto.ca/city-government/data-research-maps/research-reports/housing-and-homelessness-research-and-reports/shelter-system-flow-data/
- [40] Joanne Passaro. 2014. *The Unequal Homeless: Men on the Streets, Women in their Place*. Routledge, New York. doi:10.4324/9781315021690
- [41] Regine Paul, Emma Carmel, and Jennifer Cobbe. 2024. *Handbook on Public Policy and Artificial Intelligence*. Edward Elgar Publishing, Cheltenham, UK. doi:10.4337/9781803922171
- [42] Juan C. Perdomo, Tolani Britton, Moritz Hardt, and Rediet Abebe. 2023. Difficult Lessons on Social Prediction from Wisconsin Public Schools. arXiv:2304.06205 [cs.CY] https://arxiv.org/abs/2304.06205
- [43] Govind Persad, Monica E. Peek, and Ezekiel J. Emanuel. 2020. Fairly Prioritizing Groups for Access to COVID-19 Vaccines. *JAMA* 324, 16 (Oct. 2020), 1601–1602. doi:10.1001/jama.2020.18513
- [44] Dasha Pruss. 2023. Ghosting the Machine: Judicial Resistance to a Recidivism Risk Assessment Instrument. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (ChicAGO, IL, USA) (FAcCT '23)*. Association for Computing Machinery, New York, NY, USA, 312–323. doi:10.1145/3593013.3593999
- [45] Lisa Reutter and Heidrun Åm and. 2024. Constructing the data economy: tracing expectations of value creation in policy documents. *Critical Policy Studies* 18, 4 (2024), 639–659. doi:10.1080/19460171.2023.2300436
- [46] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3411764.3445518
- [47] Devansh Saxena, Karla Badillo-Urquiola, Pamela Wisniewski, and Shion Guha. 2021. A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021).
- [48] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2020. A Human-Centered Review of Algorithms used within the US Child Welfare System. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [49] Devansh Saxena and Shion Guha. 2024. Algorithmic Harms in Child Welfare: Uncertainties in Practice, Organization, and Street-level Decision-making. *ACM J. Responsib. Comput.* 1, 1, Article 2 (March 2024), 32 pages. doi:10.1145/3616473
- [50] Homeland Security. 2025. *DHS Playbook for Public Sector Generative Artificial Intelligence Deployment*. Retrieved Sept 2, 2025 from https://www.dhs.gov/sites/default/files/2025-01/25\_0106\_ocio\_dhs-playbook-for-public-sector-generative-artificial-intelligence-deployment-508-signed.pdf
- [51] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*. 59–68.
- [52] Ali Shirali, Rediet Abebe, and Moritz Hardt. 2024. Allocation Requires Prediction Only if Inequality Is Low. In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 235)*, Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (Eds.). PMLR, 45114–45153. https://proceedings.mlr.press/v235/shirali24a.html
- [53] Dilruba Showkat, Angela D. R. Smith, Wang Lingqing, and Alexandra To. 2023. “Who is the right homeless client?”: Values in Algorithmic Homelessness Service Provision and Machine Learning Research. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3544548.3581010

- [54] D A Stone. 1997. *Policy paradox : the art of political decision making*. W.W. Norton, New York.
- [55] Pelle Tracey and Patricia Garcia. 2024. Intermediation: Algorithmic Prioritization in Practice in Homeless Services. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW2 (Nov. 2024), 412:1–412:24.
- [56] M Walzer. 1983. *Spheres of justice : a defense of pluralism and equality*. Basic Books.
- [57] Angelina Wang, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. 2024. Against Predictive Optimization: On the Legitimacy of Decision-making Algorithms That Optimize Predictive Accuracy. *ACM J. Responsib. Comput.* 1, 1 (March 2024), 9:1–9:45. doi:10.1145/3636509
- [58] C C Yancey and M C Rourke. 2023. Emergency Department Triage.
- [59] Miri Zilka, Holli Sargeant, and Adrian Weller. 2022. Transparency, Governance and Regulation of Algorithmic Tools Deployed in the Criminal Justice System: a UK Case Study. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (Oxford, United Kingdom) (*AIES '22*). Association for Computing Machinery, New York, NY, USA, 880–889. doi:10.1145/3514094.3534200
- [60] Marta Ziosi and Dasha Pruss. 2024. Evidence of What, for Whom? The Socially Contested Role of Algorithmic Bias in a Predictive Policing Tool. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (*FAccT '24*). Association for Computing Machinery, New York, NY, USA, 1596–1608. doi:10.1145/3630106.3658991

### A Impact of relaxing equal resource allocation chances in the cutoff category in hierarchical prioritization

When we relax the assumption that everyone within the cutoff category should have an equal chance of receiving resources, within-category discrimination will also contribute to resource allocation disparities, further complicating the structural effects of hierarchical prioritization. If this condition is relaxed to allow for example, within-category weighting of subgroups  $\mathcal{S}$  then  $P_{\pi(i),s}^H$  (probability of an individual in category  $\pi(i)$  and subgroup  $s$  receiving the resource) would also be impacted by the weights assigned to the subgroup  $s$ .

### B Derivation for Equation 11

The expected number of people with attribute  $s$  who will receive the resource within the group  $i$  is:

$$\mathbb{E}[R_{\pi(i),s}] = \begin{cases} n_{\pi(i),s}, & \text{if } \sum_{j=1}^i n_{\pi(j)} \leq B, \\ n_{\pi(i),s} \left( \frac{B - \sum_{j=1}^{i-1} n_{\pi(j)}}{n_{\pi(i)}} \right), & \text{if } \sum_{j=1}^{i-1} n_{\pi(j)} < B < \sum_{j=1}^i n_{\pi(j)}, \\ 0, & \text{if } \sum_{j=1}^{i-1} n_{\pi(j)} \geq B. \end{cases}$$

Then,  $G_s^H$  is as follows where  $m$  depicts the category that marginally receives resources:

$$G_s^H = \frac{\mathbb{E}[R_s]}{N_s} = \frac{\sum_{i=1}^K n_{\pi(i),s} \cdot P_{\pi(i)}}{\sum_{i=1}^K n_{\pi(i),s}} = \frac{1}{N_s} \left[ \sum_{i=1}^{m-1} n_{\pi(i),s} + n_{\pi(m),s} \left( \frac{B - \sum_{j=1}^{m-1} n_{\pi(j)}}{n_{\pi(m)}} \right) \right]$$

Now you can differentiate the above expression.

### C Derivation for Equation 12

$\frac{d}{dB} AD^H(B)$  can be found by differentiating the below expression:

$$AD^H(B) = \left| \frac{\sum_{i=1}^{m-1} n_{\pi(i),s_1} + \frac{n_{\pi(m),s_1}}{n_{\pi(m)}} (B - \sum_{i=1}^{m-1} n_{\pi(i)})}{N_{s_1}} - \frac{\sum_{i=1}^{m-1} n_{\pi(i),s_2} + \frac{n_{\pi(m),s_2}}{n_{\pi(m)}} (B - \sum_{i=1}^{m-1} n_{\pi(i)})}{N_{s_2}} \right|$$

### D Derivation for Equation 13

Note for cutoff category  $m$ ,

$$RD^H(B) = \frac{G_{s_1}^H(B) + \epsilon}{G_{s_2}^H(B) + \epsilon} = \frac{\frac{1}{N_{s_1}} (\sum_{i=1}^{m-1} n_{\pi(i),s_1} + \frac{n_{\pi(m),s_1}}{n_{\pi(m)}} (B - \sum_{i=1}^{m-1} n_{\pi(i)})) + \epsilon}{\frac{1}{N_{s_2}} (\sum_{i=1}^{m-1} n_{\pi(i),s_2} + \frac{n_{\pi(m),s_2}}{n_{\pi(m)}} (B - \sum_{i=1}^{m-1} n_{\pi(i)})) + \epsilon}$$

Note that with a very low budget, only the first category is served and  $RD^H(B) = \frac{N_{s_2}}{N_{s_1}} \cdot \frac{n_{\pi(1),s_1}}{n_{\pi(1),s_2}}$  when we assume  $G_{s_2}^H(B) > 0$ . As  $B \rightarrow N$ , then  $RD^H(B) \rightarrow 1$ .