



# Learning Resource Allocation Policies from Observational Data with an Application to Homeless Services Delivery

Aida Rahmattalabi

University of Southern California  
USA

rahmatta@usc.edu

Kathryn Dullerud

University of Southern California  
USA

kdulleru@usc.edu

Phebe Vayanos

University of Southern California  
USA

phebe.vayanos@usc.edu

Eric Rice

University of Southern California  
USA

ericr@usc.edu

## ABSTRACT

We study the problem of learning, from observational data, fair and interpretable policies that effectively match heterogeneous individuals to scarce resources of different types. We model this problem as a multi-class multi-server queuing system where both individuals and resources arrive stochastically over time. Each individual, upon arrival, is assigned to a queue where they wait to be matched to a resource. The resources are assigned in a first come first served (FCFS) fashion according to an *eligibility structure* that encodes the resource types that serve each queue. We propose a methodology based on techniques in modern causal inference to construct the individual queues as well as learn the matching outcomes and provide a mixed-integer optimization (MIO) formulation to optimize the eligibility structure. The MIO problem maximizes policy outcome subject to wait time and fairness constraints. It is very flexible, allowing for additional linear domain constraints. We conduct extensive analyses using synthetic and real-world data. In particular, we evaluate our framework using data from the U.S. Homeless Management Information System (HMIS). We obtain wait times as low as an FCFS policy while improving the rate of exit from homelessness for underserved or vulnerable groups (7% higher for the Black individuals and 15% higher for those below 17 years old) and overall.

## CCS CONCEPTS

• Computing methodologies → Machine learning approaches.

## KEYWORDS

Fairness in AI, Causal Inference, Observational Data, Mixed-integer Optimization

### ACM Reference Format:

Aida Rahmattalabi, Phebe Vayanos, Kathryn Dullerud, and Eric Rice. 2022. Learning Resource Allocation Policies from Observational Data with an Application to Homeless Services Delivery. In *2022 ACM Conference on*

*Fairness, Accountability, and Transparency (FAccT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3531146.3533181>

## 1 INTRODUCTION

We study the problem of designing policies to effectively match heterogeneous individuals to scarce resources of different types. We consider the case where both individuals and resources arrive stochastically over time. Upon arrival, each individual is assigned to a queue where they wait to be matched to a resource. This problem arises in several public systems such as those providing social services, posing unique challenges at the intersection of efficiency and fairness. In particular, the joint characteristics of individuals and their matched resources determine the effectiveness of an allocation policy, making it crucial to match individuals with the right type of resource. Furthermore, when a resource becomes available, a decision-maker should decide whom among the individuals waiting in various queues should receive the resource which impacts the wait time of different individuals. In addition, since there are insufficient resources to meet demand, there are inherent fairness considerations for designing such policies.

We are particularly motivated by the problem of allocating housing resources among individuals experiencing homelessness. According to the U.S. Department of Housing and Urban Development (HUD), more than 580,000 people experience homelessness on a given night [20]. The Voices of Youth Count study found youth homelessness has reached a concerning prevalence level in the United States; one in 30 teens (13 to 17) and one in 10 young adults (18 to 25) experience at least one night of homelessness within a 12-month period, amounting to 4.2 million persons a year [31]. Housing interventions are widely considered as the key solution to address homelessness [33]. In the U.S., the government funds programs that assist homeless using different forms of housing interventions and services [44]. The HMIS database collects information on the provision of these services.

Unfortunately, the number of homeless individuals in the U.S. far exceeds the available resources which necessitates strategic allocation to maximize the intervention's effectiveness. Many communities have attempted to address this problem by creating coordinated community responses, typically referred to as Coordinated Entry Systems (CES). In such systems, most agencies within a community




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

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

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9352-2/22/06.

<https://doi.org/10.1145/3531146.3533181>



**Figure 1: NST-recommended resource allocation policy utilized by housing allocation agencies in the homelessness context. The policy is in the form of a resource eligibility structure. According to this figure, individuals with score eight and above qualify for PSH, score 4 to 7 are assigned to the RRH wait list and finally individuals who score below 4 are not assigned to any of the housing interventions.**

pool their housing resources in a centralized system called a Continuum of Care (CoC). A CoC is a regional or local planning body that coordinates housing and services funding—primarily from HUD—for people experiencing homelessness. Individuals in a given CoC who seek housing are first assessed for eligibility and vulnerability and those identified as having the greatest need are matched to appropriate housing resources [39]. For example, in the context of youth homelessness, the most widely adopted tool for assessing vulnerability is the Transition Age Youth-Vulnerability Index-Service Prioritization Decision Assistance Tool (TAY-VI-SPDAT): Next Step Tool (NST), which was developed by OrgCode Consulting, Corporation for Supportive Housing (CSH), Community Solutions, and Eric Rice. OrgCode claims that hundreds of CoC’s in the USA, Canada and Australia have adopted this tool [34]. After assessment, each individual receives a vulnerability score ranging from 0 to 17. One of the main challenges that CoC’s face is how to use the information about individuals to decide what housing assistance programs should be available to a particular homeless individual. In many communities, based on the recommendations provided in the NST tool documentation, individuals who score 8 to 17 are considered as “high risk” and are prioritized for resource-intensive housing programs or Permanent Supportive Housing (PSH). Those who score in the 4-7 range are typically assigned to short-term rental subsidy programs or Rapid-ReHousing (RRH) and those with score below 4 are eligible for services that meet basic needs which we refer to as Service Only (SO) [40]. Figure 1 depicts how the individuals are matched to resources according to the status-quo policy.

The aforementioned policy can be viewed as a *resource eligibility structure* as from the onset, it determines the resources an individual is eligible for. Such policies have the advantage of being interpretable, i.e., it is easy to explain why a particular allocation is made. Earlier work shows that most communities follow the policy recommendations when assigning housing [40]. However, controversy has surrounded the use of these cut scores and as of December 2020, OrgCode has called for new approaches to using the data collected by HMIS [35]. There is also an overwhelming desire on the part of HUD to design systematic and data-driven housing policies, including the design of the cut scores and the queues that they induce [44]. Currently, the cut scores are not tied to the treatment effect of interventions or the relative arrival rate

of individuals and resources in the respective queues. This is problematic as it is not evidently clear that assigning high-scoring and mid-scoring individuals to particular housing interventions, such as PSH or RRH, actually increases their chances of becoming stably housed. Additionally, there may not be enough resources to satisfy the needs of all individuals matched to a particular resource, resulting in long wait times. Prolonged homelessness may in turn increase the chances of exposure to violence, substance use, etc., or individuals dropping out of the system.

In particular, OrgCode and others have called for a new equity focus to how vulnerability tools are linked to housing allocation [29, 35]. Despite recent efforts to understand and mitigate disparities in homelessness, current system suffers from a significant gap in the prevalence of homelessness across different groups. For example, studies show that most racial minority groups experience homelessness at higher rates than Whites [18]. Also, recent work has revealed that PSH outcomes are worse for Black clients in Los Angeles [29] and based on the same HMIS data used in present study, Black, Latinx, and LGBQ youth have been shown to experience worse housing outcomes [22]. Addressing these disparities requires an understanding of the distribution of the individuals vulnerable to homelessness, the heterogeneity in the treatment affect and the associations with protected attributes such as race, gender, or age.

In this work, we build on the literature on causal inference and queuing theory and propose a methodology that uses historical data about the waitlisted individuals and their allocated resources to optimize resource allocation policies. We make the following contributions:

- We model the policy optimization problem as a multi-class multi-server queuing system between heterogeneous individuals and resources that arrive over time. We extend the literature on queuing theory by proposing a data-driven methodology to construct the model from observational data. Specifically, we use tools from modern causal inference to learn the treatment effect of the interventions from data and construct the queues by grouping individuals that have similar average treatment effects.
- We propose interpretable policies that take the form of a resource eligibility structure, encoding the resource types that serve each queue. We provide an MIO formulation to optimize the eligibility structure that incorporates flexibly defined fairness considerations or other linear domain-specific constraints. The MIO maximizes the policy effectiveness and guarantees minimum wait time.
- Using HMIS data, we conduct a case study to demonstrate the effectiveness of our approach. Our results indicate superior performance along policy effectiveness, fairness and wait time. Precisely, we are able to obtain wait time as low as a fully FCFS policy while improving the rate of exit from homelessness for traditionally underserved or vulnerable groups (7% for the Black individuals and 15% higher for youth below 17 years old) and overall.

The remainder of this paper is organized as follows. In Section 2, we review the related literature. In Section 3, we introduce the policy optimization problem. In Section 4, we propose our data-driven

methodology for solving the policy optimization problem. Finally, we summarize our numerical experiments and present a case study using HMIS data on youth experiencing homelessness in Section 5. Proofs and detailed numerical results are provided in the Appendix.

## 2 LITERATURE REVIEW

This work is related to several streams of literature which we review. Specifically, we cover queuing theory as the basis of our modelling framework. We also position our methodology within the literature on data-driven policy optimization and causal inference. We conclude by highlighting recent works on fairness in resource allocation.

A large number of scarce resource allocation problems give rise to one-sided queuing models. In these models, resources are allocated upon arrival, whereas individuals queue before being matched. Examples are organ matching [7] and public housing assignment [4, 24]. One stream of literature studies dynamic matching policies to find asymptotically optimal scheduling policies under conventional heavy traffic conditions [5, 28]. Another stream focuses on the system behavior under FCFS service discipline aiming to identify conditions that ensure the stability of the queuing system and characterize the steady-state matching flow rates, i.e., the average rate of individuals of a given queue (or customer class) that are served by a particular resource (server) [10, 17]. These works only focus on minimizing delay and do not explicitly model the heterogeneous service value among the customers. Recently, [15] studied one-sided queuing system where resources are allocated to the customer with the highest score (or index), which is the sum of the customer's waiting score and matching score. The authors derive a closed-form index that optimizes the steady-state performance subject to specific fairness considerations. Their proposed fairness metric measures the variance in the likelihood of getting service before abandoning the queue. Contrary to their model, we consider FCFS policies subject to resource eligibility structures which we optimize over. Our model is based on the policies currently being implemented for housing allocation among homeless individuals that target resources to heterogeneous individuals by explicitly imposing an eligibility structure. Our policies are interpretable by design as upon arrival the resources that an individual is eligible for is known, making it easy to explain why a certain allocation has or has not been made. Further, the proposed model allows for a more general class of fairness requirements commonly used in practice including fairness in *allocation* and *outcome*. It is noteworthy that our model is different from common allocation models in the public housing setting where targeting of the resources can be considered as implicit, i.e., individuals with different levels of need make different choices about where to apply and what to accept themselves [3, 4].

Our approach builds upon [2], in which the authors study the problem of designing a matching topology between customer classes and servers under a FCFS service discipline. They focus on finding matching topologies that minimize the customers' waiting time and maximize matching rewards obtained by pairing customers and servers. The authors characterize the average steady-state wait time across all customer classes in terms of the structure of the matching model, under heavy-traffic condition. They propose a quadratic

program (QP) to compute the steady-state matching flows between customers and servers and prove the conditions under which the approximation is exact. We build on the theoretical results in [2] to design resource eligibility structures that match heterogeneous individuals and resources in the homelessness setting. Contrary to the model in [2], we do not assume that the queues or the matching rewards are given a priori. Instead, we propose to use observational data from historical policy to learn an appropriate grouping of individuals into distinct queues, estimate the matching rewards, and evaluate the resulting policies.

Another stream of literature focuses on designing data-driven policies, where fairness considerations have also received significant attention due to implicit or explicit biases that models or the data may exhibit [9, 14, 25, 36]. In [9], the authors propose a data-driven model for learning scoring policies for kidney allocation that matches organs at their time of procurement to available patients. Their approach satisfies linear fairness constraints approximately and does not provide any guarantees for wait time. In addition, they take as input a model for the matching rewards (i.e., life years from transplant) to optimize the scoring policy. In [6], the authors propose a data-driven mixed integer program with linear fairness constraints to solve a similar resource allocation which provides an exact, rather than an approximate, formulation. They also give an approximate solution to achieve faster run-time. We consider a class policies in the form of matching topologies that is different from scoring rules and is more closely related to the policies implemented in practice. Such policies offer more interpretability as individuals know what resources they are eligible for from the onset. Several works have considered interpretable functional forms in policy design. For example, in [8, 23], the authors consider decision trees and develop techniques to obtain optimal trees from observational data. Their approach is purely data-driven and do not allow for explicit modelling of the arrival of resources, individuals which impact wait time. In the homelessness setting, our work is closely related to [27] which proposes a resource allocation mechanism to match homeless households to resources based on the probability of system re-entry. In this work, the authors provide a static formulation of the problem which requires frequent re-optimization and does not take the waiting time into account. In [32], the authors propose a fairness criterion that prioritizes those who benefit the most from a resource, as opposed to those who are the neediest and study the price of fairness under different fairness definitions. Similar to [27], their formulation is static and does not yield a policy to allocate resources in dynamic environments.

## 3 HOUSING ALLOCATION AS A QUEUING SYSTEM

### 3.1 Preliminaries

We model the resource allocation system as an infinite stream of heterogeneous individuals and resources that arrive over time. Each individual is characterized by a (random) feature vector  $X \in \mathcal{X} \subseteq \mathbb{R}^n$  and receives an intervention  $R$  from a finite set of treatments indexed in the set  $\mathcal{R}$ . We note that  $\mathcal{R}$  may include "no intervention" or minimal interventions such as SO in the housing allocation setting. Using the potential outcomes framework [42], each individual has

a vector of potential outcomes  $Y(r) \in \mathcal{Y} \subseteq \mathbb{R} \forall r \in \mathcal{R}$ , where  $Y(r)$  is an individual's outcome when matched to resource  $r$ .

We assume having access to  $N$  historical observations  $\mathcal{D} := \{(X_i, R_i, Y_i)\}_{i=1}^N$ , generated by the deployed policy, where  $X_i \in \mathcal{X}$  denotes the feature vector of the  $i$ th observation,  $R_i \in \mathcal{R}$  is the resource assigned to it and  $Y_i = Y_i(R_i)$  is the observed outcome, i.e., the outcome under the resource received. A (stochastic) policy  $\pi(r|\mathbf{x}) : \mathcal{X} \times \mathcal{R} \rightarrow [0, 1]$  maps features  $\mathbf{x}$  to the probability of receiving resource  $r$ . We define the value of a policy as the expected outcome when the policy is implemented, i.e.,  $V(\pi) := \mathbb{E}[\sum_{r \in \mathcal{R}} \pi(r|\mathbf{x}) Y(r)]$ . A major challenge in evaluating and optimizing policies is that we cannot observe the counterfactual outcomes  $Y_i(r), r \in \mathcal{R}, r \neq R_i$  of resources that were not received by data point  $i$ . Hence, we need to make further assumptions to identify policy values from historical data. In Section 4, we elaborate on these assumptions and propose our methodology for evaluating and optimizing policies from data.

We model the system as a multi-class multi-server (MCMS) queuing system where a set of resources  $\mathcal{R}$  serve a finite set of individual queues indexed in the set  $\mathcal{Q}$ . Upon arrival, individuals are assigned to different queues based on their feature vector. For example, in the housing allocation setting and according to the recommended policy the assignment is based on the vulnerability score. We use  $p : \mathcal{X} \rightarrow \mathcal{Q}$  to denote the function that maps the feature vector to a queue that the individual will join. We refer to  $p$  as the *partitioning function* (as it partitions  $\mathcal{X}$  and assigns each subset to a queue) and note that it is unknown a priori. In this work, we consider partitioning functions in the form of a binary trees similar to classification trees, due to their interpretability [6]. We assume that individuals arrive according to stationary and independent Poisson processes and that inter-arrival time of resources follows an exponential distribution. These are common assumptions in queuing theory for modeling arrivals, however, they may not fully hold in practice in which case, it is possible to use re-optimization to adapt to the changing environment. We use  $\lambda := (\lambda_1, \dots, \lambda_{|\mathcal{Q}|})$  and  $\mu := (\mu_1, \dots, \mu_{|\mathcal{R}|})$  to denote the vector of arrival rates of individuals and resources, respectively. We define  $\lambda_Q := \sum_{q \in \mathcal{Q}} \lambda_q$  and  $\mu_{\mathcal{R}} := \sum_{r \in \mathcal{R}} \mu_r$  as the cumulative arrival rates of individuals and resources, respectively. Without loss of generality, we assume that  $\lambda_q > 0 \forall q \in \mathcal{Q}$  and  $\mu_r > 0 \forall r \in \mathcal{R}$ .

### 3.2 Matching Policy

Once a new resource becomes available, it is allocated according to a resource eligibility structure that determines what queues are served by any particular resource. The resource eligibility structure can be represented as a matching topology  $M := [M_{qr}] \in \{0, 1\}^{|\mathcal{Q}| \times |\mathcal{R}|}$ , where  $M_{qr} = 1$  indicates that individuals in queue  $q$  is eligible for resource  $r$ . Resources are assigned to queues in an FCFS fashion subject to matching topology  $M$ . For a partitioning function  $p$  and matching topology  $M$ , we denote the allocation policy by  $\pi_{p,M}(r|\mathbf{x})$ . We concern ourselves with the long-term steady state of the system. Proposition 1 gives the necessary and sufficient conditions to arrive at a steady-state.

**Proposition 1** (Adan and Weiss [1], Theorem 2.1). *Given the MCMS system defined through  $(\mathcal{Q}, \mathcal{R}, \lambda, \mu, M)$ , under the FCFS service discipline matching  $M$  admits a steady state if and only if the following*

*condition is satisfied:*

$$\mu_{\mathcal{R}} - \sum_{r \in \mathcal{R}} \sum_{q \in \mathcal{Q}_M(\mathbf{M})} \lambda_q > 0 \quad \forall \mathcal{R} \subseteq \mathcal{R}.$$

*The left-hand side is the cumulative arrival rate of resources in  $\mathcal{R}$  in excess of the cumulative arrival rate of all the queues in  $\mathcal{Q}_M$ , where  $\mathcal{Q}_M$  is the set of queues that are only eligible for resources in  $\mathcal{R}$ , i.e.,  $\mathcal{Q}_M = \{q \in \mathcal{Q} : \sum_{r \in \mathcal{R} \setminus \mathcal{R}} M_{qr} = 0\}$ .*

We define the set of *admissible matching topologies* as those that satisfy the inequality in Proposition 1. In the housing allocation problem, we assume that SO resources are abundant, i.e.,  $\mu_{\mathcal{R}} - \lambda_Q > 0$ . In other settings, it is possible to create an auxiliary resource queue corresponding to no intervention. The abundance assumption ensures that there exists at least one admissible matching: the fully connected matching topology  $M_{qr} = 1 \forall q \in \mathcal{Q}, r \in \mathcal{R}$  and is necessary in order to avoid overloaded queues. In practice housing resources are strictly preferred. As a result, we propose to study the system under the so-called *heavy traffic* regime, where the system is loaded very close to its capacity and we assume that the system utilization parameter  $\rho := \mu_{\mathcal{R}}/\lambda_Q$  approaches 1, i.e.,  $\rho \rightarrow 1$ . In general, we assume that  $\lambda$  and  $\mu$  are such that  $\lambda_Q = \rho \mu_{\mathcal{R}}$ . This assumption will additionally make the analytical study of the matching system more tractable. In particular, in [2], the authors propose a quadratic program to approximate the exact steady-state flows of the stochastic FCFS matching system under heavy traffic conditions. They enforce the steady-state flows in an optimization model to find the optimal matching topology using KKT optimality conditions. We adopt the same set of constraints in Section 3.4 where we present the optimization formulation. We let  $F := [F_{qr}] \in \mathbb{R}_+^{|\mathcal{Q}| \times |\mathcal{R}|}$  denote the steady-state flow, where  $M_{qr} = 0 \Rightarrow F_{qr} = 0$ . Given a partitioning function  $p$ , the policy associated with a matching topology  $M$  is equal to  $\pi_{p,M}(r|\mathbf{x}) = F_{qr}/\sum_{r \in \mathcal{R}} F_{qr} = F_{qr}/\lambda_q$ , in which  $q = p(\mathbf{x})$  and the second inequality follows from the flow balance constraints. In Proposition 2 we show how the policy value can be written using the matching model parameters and treatment effect of different interventions. We define the conditional average treatment effect (CATE) of resource  $r$  and queue  $q$  as  $\tau_{qr} := \mathbb{E}[Y(r) - Y(1)|P(X) = q] \forall r \in \mathcal{R}, q \in \mathcal{Q}$ , in which  $r = 1$  is the baseline intervention. In many applications, the baseline intervention corresponds to “no-intervention” (also referred to as the control group). In the housing allocation context, we set  $r = 1$  to be the SO intervention.

**Proposition 2.** *Given a partitioning function  $p$ , an MCMS model  $(\mathcal{Q}, \mathcal{R}, \lambda, \mu, M)$ , and the steady-state FCFS flow  $F$  under FCFS discipline, the value of the induced policy  $\pi_{p,M}$  is equal to:*

$$V(\pi_{p,M}) = \frac{1}{\lambda_Q} \sum_{q \in \mathcal{Q}} \sum_{r \in \mathcal{R}} F_{qr} \tau_{qr} + C,$$

*where  $C$  is a constant that depends on the expected outcome under the baseline intervention.*

### 3.3 Policy Optimization

We now introduce the policy optimization problem under the assumption that the joint distribution of  $X, Y(r), r \in \mathcal{R}$  as well as the partitioning function  $p$  is known. In Section 4, we propose a

methodology to construct  $p$ . Furthermore, we describe how we can use historical data to optimize new policies.

$$\mathcal{P}(p) := \max_{M \in \mathcal{M}} V(\pi_{p,M}). \quad (1)$$

In this problem,  $\mathcal{M}$  is the set of feasible matching topologies which incorporates constraints that impose steady-state flow, fairness and minimum wait time constraints.

**Fairness.** In this work, we focus on group-based notions of fairness which have been widely studied in recent years in various data-driven decision making settings [6, 9, 32, 37]. Formally, we let  $G$  be a random variable describing the group that an individual belongs to, taking values in  $\mathcal{G}$ . For example,  $G$  can correspond to protected features such as race, gender or age. It is also possible to define fairness with respect to other features, such as vulnerability score in the housing allocation setting. We give several examples to which our framework applies.

**Example 1** (Maximin Fairness in Allocation). *Motivated by Rawls theory of social justice [38], maximin fairness aims to help the worst-off group as much as possible. Formally, the fairness constraints can be written as*

$$\sum_{q \in Q_g} F_{qr} \geq w \quad \forall g \in \mathcal{G}, r \in \mathcal{R},$$

where  $w$  is the minimum acceptable flow across groups and  $Q_g \subseteq Q$  is a subset of queues whose individuals belong to  $G = g$ . If queues contain individuals with different values of  $g$ , one should separate them by creating multiple queues with unique  $g$ . By increasing the parameter  $w$ , one is imposing more strict fairness requirements. This parameter can be used to control the trade-off between fairness and policy value. It can also be set to the highest value for which the constraint is feasible.

**Example 2** (Group-based Parity in Allocation). *Parity-based fairness notions strive for equal outcomes across groups.*

$$\left| \sum_{q \in Q_g} F_{qr} - \sum_{q \in Q_{g'}} F_{qr} \right| \leq \epsilon \quad \forall g, g' \in \mathcal{G}, r \in \mathcal{R}.$$

In words, for every resource the difference between the cumulative flow between any pair of groups should be at most  $\epsilon$ , where  $\epsilon$  can be used to control the trade-off between fairness and policy value.

**Example 3** (Maximin Fairness in Outcome). *For every group, the policy value should be at least  $w$ .*

$$\frac{1}{\lambda_{Q_g}} \sum_{q \in Q_g} \sum_{r \in \mathcal{R}} F_{qr} \tau_{qr} \geq w \quad \forall g \in \mathcal{G}.$$

**Example 4** (Group-based Parity in Outcome). *The difference between the policy value for any pair of groups is at most  $\epsilon$ .*

$$\left| \frac{1}{\lambda_{Q_g}} \sum_{q \in Q_g} \sum_{r \in \mathcal{R}} F_{qr} \tau_{qr} - \frac{1}{\lambda_{Q_{g'}}} \sum_{q \in Q_{g'}} \sum_{r \in \mathcal{R}} F_{qr} \tau_{qr} \right| \leq \epsilon \quad \forall g, g' \in \mathcal{G}.$$

In the experiments, we focus on fairness in outcome due to treatment effect heterogeneity. In other words, it is important to match individuals with the right type of resource, rather than ensuring all groups have the same chance of receiving any particular resource. Further, we adopt maximin fairness which guarantees Pareto optimal policies [36].

**Wait Time.** Average wait time is dependent on the structure of the matching topology. For example, minimum average wait time is attainable in a fully FCFS policy where  $M_{qr} = 1 \forall q \in Q, r \in \mathcal{R}$ . In [2], the authors characterize the general structural properties that impact average wait time. In particular, they show that under the heavy traffic condition, a matching system can be partitioned into a collection of complete resource pooling (CRP) subsystems that operate “almost” independently of each other. A key property of this partitioning is that individuals that belong to the same CRP component experience the same average steady-state wait time. Furthermore, the average wait time is tied to the number of CRPs of a matching topology, where a single CRP achieves minimum average wait time. In [2], the authors introduce necessary and sufficient constraints to ensure that the matching topology  $M$  induces a single CRP component. We adopt these constraints in order to achieve minimum wait time which we discuss next. Next, we present the final optimization formulation by incorporating both wait time and fairness considerations in Problem (1).

### 3.4 Optimization Formulation

Suppose the joint distribution of  $X, Y(r) \forall r \in \mathcal{R}$  is known. Given a partitioning function  $p$ , problem (1) can be solved via the MIO:

$$\max \sum_{q \in Q} \sum_{r \in \mathcal{R}} \tau_{qr} F_{qr} \quad (2a)$$

$$\text{s.t. } F_{qr}, v_{qr} \in \mathbb{R}_+, \gamma_r, \theta_q \in \mathbb{R} \quad \forall q \in Q, r \in \mathcal{R} \quad (2b)$$

$$M_{qr}, z_{qr} \in \{0, 1\} \quad \forall q \in Q, r \in \mathcal{R} \quad (2c)$$

$$g_{qr}^{(k)} \in \mathbb{R}_+ \quad \forall q, k \in Q, r \in \mathcal{R} \quad (2d)$$

$$\sum_{q \in Q} F_{qr} = \mu_r \quad \forall r \in \mathcal{R} \quad (2e)$$

$$\sum_{r \in \mathcal{R}} F_{qr} = \lambda_q \quad \forall q \in Q \quad (2f)$$

$$F_{qr} \leq \lambda_q \mu_r (\theta_q + \gamma_r + v_{qr}) + Z(1 - m_{qr}) \quad \forall q \in Q, r \in \mathcal{R} \quad (2g)$$

$$F_{qr} \geq \lambda_q \mu_r (\theta_q + \gamma_r + v_{qr}) - Z(1 - m_{qr}) \quad \forall q \in Q, r \in \mathcal{R} \quad (2h)$$

$$F_{qr} \leq Z m_{qr} \quad \forall q \in Q, r \in \mathcal{R} \quad (2i)$$

$$F_{qr} \leq Z z_{qr} \quad \forall q \in Q, r \in \mathcal{R} \quad (2j)$$

$$v_{qr} \leq (|Q| + |\mathcal{R}| + 1)W(1 - z_{qr}) \quad \forall q \in Q, r \in \mathcal{R} \quad (2k)$$

$$\sum_{q \in C} g_{qr}^{(k)} = \mu_r \quad \forall r \in \mathcal{R}, k \in Q \quad (2l)$$

$$\sum_{r \in \mathcal{R}} h_{qr}^{(k)} = \lambda_q - \frac{\delta}{|Q| - 1} \quad \forall q \in Q \setminus \{k\}, k \in Q \quad (2m)$$

$$\sum_{r \in \mathcal{R}} g_{qr}^{(k)} \leq Z m_{qr} \quad \forall q, k \in Q, r \in \mathcal{R} \quad (2n)$$

$$\sum_{r \in \mathcal{R}} g_{qr}^{(k)} = \lambda_k + \delta \quad \forall q, k \in Q \quad (2o)$$

$$\frac{1}{\lambda_{Q_g}} \sum_{q \in Q_g} \sum_{r \in \mathcal{R}} f_{qr} \tau_{qr} + C \geq w \quad \forall g \in \mathcal{G}, F \in \mathcal{F}. \quad (2p)$$

In this formulation,  $\delta := \left( \prod_{q \in Q} v_q \prod_{r \in \mathcal{R}} v_r \right)^{-1}$  and  $\frac{w_q}{v_q} = \lambda_q$ ,  $\frac{w_r}{v_r} = \mu_r$  are rational number representations. Constants  $W$  and  $Z$  are defined as follows:  $W := 1/2 \max\{\max_{q \in Q} 1/\lambda_q, \max_{r \in \mathcal{R}} 1/\mu_r\}$ , and  $Z := \lambda^{(\max)} \mu^{(\max)} \left( \sum_{q \in Q} 1/\lambda_q + \sum_{r \in \mathcal{R}} 1/\mu_r + (|Q| + |\mathcal{R}| + 1)^2 W \right)$ , where  $\lambda^{(\max)}$  and  $\mu^{(\max)}$  are the maximum arrival rate across demand (arriving individuals) and resource queues, respectively. Constraints (2e) and (2f) are the flow balance constraints. Constants  $W, Z$  ensure that constraints (2g)-(2k) impose the KKT conditions of the quadratic program that approximates steady-state-flow for a matching topology  $\mathcal{M}$ . Constraints (2l)-(2o) enforce a single CRP component to ensure minimum wait time. Finally, constraint (2p) collects the fairness constraints where we can use any of the aforementioned examples. In order to solve problem (2), we need to estimate  $\tau_{qr}$  and  $\lambda$  which depend of  $p$ , as well as  $\mu$ . Once the queues are fixed, estimating  $\lambda$  values are straightforward. In addition,  $\mu$  can be easily estimated from the historical provision of resources. In the next section, we discuss how to construct the queues and simultaneously estimate the  $\tau_{qr}$  values.

## 4 SOLUTION APPROACH

We first partition  $\mathcal{X}$  and then estimate CATE in each subset of the partition. We propose to use causal trees to achieve both tasks simultaneously [45]. Causal trees estimate CATE of binary interventions by partitioning the feature space into sub-populations that differ in the magnitude of their treatment effects. The method is based on regression trees, modified to estimate the goodness-of-fit of treatment effects. A key aspect of using causal trees for partitioning is that the cut points on features are such that the treatment effect variance within each leaf node is minimized. In other words, individuals who are similar in the treatment effect are grouped together in a leaf node. This results in queues that are tied to the treatment effect of resources which will result in improved policy value (see Section 5).

### 4.1 Assumptions

Causal trees rely on several key assumptions which are standard in causal inference for treatment effect estimation [21]. These assumptions are usually discussed for the case of binary treatments. Below, we provide a generalized form of the assumptions for multiple treatments.


**Assumption 1** (Stable Unit Treatment Value Assumption (SUTVA)). *The treatment that one unit (individual) receives does not change the potential outcomes of other units.*

**Assumption 2** (Consistency). *The observed outcome agrees with the potential outcome under the treatment received.*

The implication of this assumption is that there are no different forms of each treatment which lead to different potential outcomes. In the housing allocation setting, this requires that there is only one version of PSH, RRH and SO.

**Assumption 3** (Positivity). *For all feature values, the probability of receiving any form of treatment is strictly positive, i.e.,*

$$\mathbb{P}(R = r | X = \mathbf{x}) > 0 \quad \forall r \in \mathcal{R} \quad \mathbf{x} \in \mathcal{X}.$$



**Figure 2: Example partitioning by sample causal trees for PSH and RRH interventions.**

The positivity assumption states that any individual should have a positive probability of receiving any treatment. Otherwise, there is no information about the distribution of outcome under some treatments and we will not be able to make inferences about it. In Section 5, we discuss the implications of this assumption in the context of HMIS data.

**Assumption 4** (Conditional Exchangeability). *Individuals receiving a treatment should be considered exchangeable, with respect to potential outcomes  $Y(r)$ ,  $r \in \mathcal{R}$ , with those not receiving it and vice versa. Mathematically,*

$$Y(1), \dots, Y(|\mathcal{R}|) \perp R | X = \mathbf{x} \quad \forall \mathbf{x} \in \mathcal{X}.$$

Conditional exchangeability means that there are no unmeasured confounders that are a common cause of both treatment and outcomes. If unmeasured confounders exist, it is impossible to accurately estimate the causal effects. In observational settings, a decision-maker only relies on passive observations. As a result, in order to increase the plausibility of this assumption, researchers typically include as many features as possible in  $X$  to ensure that as many confounders as possible between treatment and outcome are accounted for. In the housing allocation setting, the HMIS data contains a rich set of features (54 features) associated with different risk factors for homelessness which we use in order to estimate the treatment effects. In Section 6, we discuss the consequences of violating the above assumptions.

### 4.2 Building the Partitioning Function

Next, we describe our approach for estimating CATE. We first consider a simple case with binary treatments, i.e.,  $|\mathcal{R}| = 2$  as causal trees work primarily for binary treatments. After training the causal tree using the data on a pair of treatments, the leaves induce a partition on the feature space  $\mathcal{X}$ . Hence, we can view the causal tree as the partitioning function  $p$  where each individual is uniquely mapped to a leaf node, i.e., a queue.

Extending to the case of  $|\mathcal{R}| > 2$  is non-trivial. Assuming  $r = 1$  is the baseline intervention, we construct  $|\mathcal{R}| - 1$  separate causal trees to estimate CATE for  $r \in \mathcal{R} \setminus \{1\}$ . We denote the resulting causal trees or partitioning functions  $p_r : \mathcal{X} \rightarrow Q \quad \forall r \in \mathcal{R} \setminus \{1\}$ . We define  $\mathcal{X}_r(q) = \{\mathbf{x} \in \mathcal{X} : p_r(\mathbf{x}) = q\} \quad \forall r \in \mathcal{R} \setminus \{1\}, q \in Q$  as the set of all individuals who belong to queue  $q$  according to partitioning function  $p_r$ . Also, let  $q_r = p_r(\mathbf{x})$ . In order to aggregate the individual partitioning functions to obtain a unified partition on  $\mathcal{X}$ , we consider the intersections of  $\mathcal{X}_r(q)$  created by each tree. We define

subsets  $\mathcal{X}(q_1, \dots, q_{|\mathcal{R}|-1}) = \bigcap_{r=1}^{|\mathcal{R}|-1} \mathcal{X}_r(q_r)$  for all combinations of  $q_r \in \mathcal{Q}$ . We can view  $\mathcal{X}(q_1, \dots, q_{|\mathcal{R}|-1})$  as a new (finer) partition on  $\mathcal{X}$ . We illustrate with an example using the housing allocation setting. Suppose we have constructed two causal trees for PSH and RRH according to Figure 2 such that PSH tree splits the vulnerability score into intervals of  $[0, 6]$ ,  $[7, 10]$ ,  $[11, 17]$  and RRH creates  $[0, 8]$ ,  $[9, 17]$  subsets. According to our procedure, the final queues are constructed using the intersection of these subsets. In other words, we obtain  $[0, 6]$ ,  $[7, 8]$ ,  $[9, 10]$ ,  $[11, 17]$  which corresponds to four queues. Building queues as the intersection of the tree subsets may result in a large number of queues. However, even though the size of the optimization problem is quadratic in the number of queues, in practice the problem solves relatively quickly (in seconds) for queue number up to 30 which is far more than what is typically seen in practice. In addition, the granularity of the partition, and subsequently the number of queues, can be controlled through the tree depth or the minimum allowable number of data points in each leaf, both of which are adjustable parameters in causal trees. It is noteworthy that while fewer queues result in models that are easily interpretable, more queues allow the decision-makers to leverage the heterogeneity in treatment effect to target the resources to the right individuals, hence achieving higher-valued policies. We explore this trade-off in the experimental result section.

Finally, in order to estimate  $\tau_{qr}$ , we should avoid using the estimates from each individual tree. The reason is that each tree estimates  $\mathbb{E}[Y(r) - Y(1)|p(X) = q, R \in \{1, r\}] \forall r \in \mathcal{R} \setminus \{1\}$ . That is, a subset of the data associated with a pair of treatments is used to build each tree. Therefore,  $\tau_{qr}$  values are not generalizable to the entire population and need to be re-evaluated over all data points that belong to a subset. We adopt Doubly Robust estimator (DR) for this task. Proposed in [16], DR combines an outcome regression with a model for the treatment assignment (propensity score) to estimate treatment effects. DR is an unbiased estimate of treatment effects, if at least one of the two models are correctly specified. Hence, it has a higher chance of reliable inference. CATE estimates  $\hat{\tau}_{qr}$  are provided below.

$$\hat{\tau}_{qr} = T(r) - T(1) \quad r \in \mathcal{R},$$

where  $T(r) = \frac{1}{|I_q|} \sum_{i \in I_q} \left( \hat{y}(X_i, r) + (Y_i - \hat{y}(X_i, R_i)) \frac{\mathbb{I}(R_i=r)}{\bar{\pi}(R_i|X_i)} \right)$  and  $I_q := \{i \in \{1, \dots, N\} : p(X_i) = q\}$  is the set of indices in the historical data that belongs to  $q$ . Further,  $\hat{y}$  and  $\bar{\pi}$  are the outcome and historical policy (i.e., propensity score) models, respectively. According to the above expression, all resources are compared to the baseline intervention, hence  $\hat{\tau}_{qr} = 0$  for  $r = 1$ .

We end this section by discussing a practical consideration which is a desire to design policies that depend on low-dimensional features, such as risk scores. In cases that we only use risk scores, not the full feature vector, it is critical that they satisfy the causal assumptions. We provide a risk score formulation that satisfies this requirement.

**Proposition 3.** *We define risk score functions as  $S_r(\mathbf{x}) = \mathbb{P}[Y(r) = 1|X = \mathbf{x}] \forall r \in \mathcal{R}$ . Suppose  $S \in \mathcal{S}$  is a (random) vector of risk scores. Also, let  $Y = (Y(1), \dots, Y(|\mathcal{R}|))$  be the vector of potential outcomes. The following statements hold for all  $\forall \mathbf{x} \in \mathcal{X}, s \in \mathcal{S}$ :*

- (1)  $Y \perp R|X \Rightarrow Y \perp R|S$ .
- (2)  $\mathbb{P}(\mathbb{P}(R = r|X = \mathbf{x}) > 0) = 1 \Rightarrow \mathbb{P}(\mathbb{P}(R = r|S = \mathbf{s}) > 0) = 1$ .

Under causal assumptions,  $S_r(\mathbf{x}) = \mathbb{P}(Y(r) = 1|X = \mathbf{x}, R = r) = \mathbb{P}(Y = 1|X = \mathbf{x}, R = r)$ , which relies on observed data, rather than counterfactuals. According to Proposition 3, as in general individuals respond differently to various treatments, one risk score per resource may be required in order to summarize the information of  $X$ . Alternatively, one can utilize the entire set of features in the causal tree and to learn the propensity and outcome models used in the treatment effect estimators.

## 5 COMPUTATIONAL RESULTS

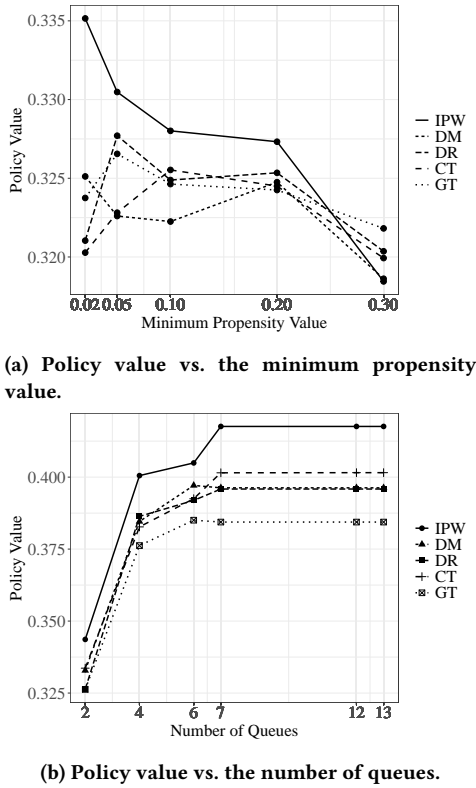
We conduct two sets of experiments to study the performance of our approach to design resource allocation policies: (i) synthetic experiments where the treatment and potential outcomes are generated according to a known model; (ii) experiments on the housing allocation system based on HMIS data for youth experiencing homelessness. We use the causal tree implementation in the *grf* package in R. We control the partition granularity by changing the *minimum node size* parameter which is minimum number of observations in each tree leaf. We evaluate policies using three estimators from the causal inference literature [16]: Inverse Propensity Weighting (IPW) which corrects the mismatch between the historical policy and new policy by re-weighting the data points with their propensity values, Direct Method (DM) which uses regression models to estimate the unobserved outcomes, and DR. In addition, we include objective value of Problem (2) obtained by matching flow and CATE estimates (CT). When models of outcome and propensity are correctly specified, the above estimators are all unbiased [16].

### 5.1 Synthetic Experiments

We generate synthetic potential outcomes and resource assignments in the HMIS data collected between 2015 and 2017 from 16 communities across the United States [11]. We use the following setting using vulnerability score  $S$  (unless mentioned otherwise):  $\bar{\pi}(\text{SO}|S > 0.2) = 0.3$ ,  $\bar{\pi}(\text{SO}|0.0 < S \leq 0.2) = 0.3$  and  $\bar{\pi}(\text{SO}|S \leq 0.0) = 0.3$ . Additionally,  $\bar{\pi}(\text{RRH}|S > 0.2) = 0.2$ ,  $\bar{\pi}(\text{RRH}|0.0 < S \leq 0.2) = 0.4$  and  $\bar{\pi}(\text{RRH}|S \leq 0.0) = 0.3$  and finally,  $\bar{\pi}(\text{PSH}|S > 0.2) = 0.5$ ,  $\bar{\pi}(\text{PSH}|0.0 < S \leq 0.2) = 0.3$  and  $\bar{\pi}(\text{PSH}|S \leq 0.0) = 0.4$ . The potential outcomes are sampled from binomial distributions with probabilities that depend on  $S$ . For PSH, we use  $\mathbb{E}[Y(\text{PSH})|S \leq 0.3] = 0.6$ ,  $\mathbb{E}[Y(\text{PSH})|0.3 < S \leq 0.5] = 0.2$  and  $\mathbb{E}[Y(\text{PSH})|0.5 < S] = 0.6$ . For RRH,  $\mathbb{E}[Y(\text{RRH})|S \leq 0.2] = 0.2$ ,  $\mathbb{E}[Y(\text{RRH})|0.2 < S \leq 0.7] = 0.6$  and  $\mathbb{E}[Y(\text{RRH})|0.7 < S] = 0.2$ . Finally,  $\mathbb{E}[Y(\text{SO})] = 0$ . We evaluate policies obtained by solving Problem (2). We use decision trees for outcome and propensity score models.

One of the goals of the synthetic experiments is to compare different estimators in a setting where we observe the potential outcomes. Specifically, we study the performance of the estimators for policy evaluation when propensity values are varied. We generate different datasets by changing the propensity values  $\bar{\pi}(\text{PSH}|0.0 < S \leq 0.2) = \alpha$  and  $\bar{\pi}(\text{RRH}|0.0 < S \leq 0.2) = 0.7 - \alpha$  for  $\alpha \in \{0.02, 0.05, 0.1, 0.2, 0.3\}$  and obtain the optimal policy for each dataset. Figure 3a shows optimal policy values according to different estimators. We observe that across the  $x$ -axis range, DR, DM and CT result in similar estimates which also agrees with the ground truth (GT). However, when the minimum propensity score is small ( $< 0.05$ ), IPW diverges from GT. This is consistent with





**Figure 3: Synthetic data. Each line corresponds to a different estimator.**

other findings in the literature suggesting that when propensities are too close to 0 or 1, non-parametric estimators tend to have higher variance and converge at a slower rate (with the number of data points) [26].

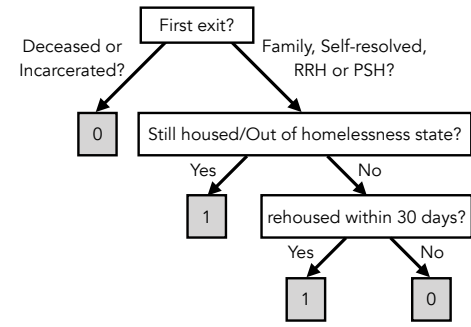
Next, we investigate the effect of treatment heterogeneity on the value of the optimal policy. In particular, we study how much the granularity of partitions, or the number of queues, impacts the policy value. Figure 3b summarizes the results. When the number of queues is equal to 1, the optimal policy is at its minimum value. In this case, the policy corresponds to an FCFS policy as individuals queue in a single line and are prioritized according to their arrival times. The optimal policy value gradually increases ( $\sim 25\%$  according to GT) as the number of queues increases until it flattens. This suggests that by increasing the number of queues, we can leverage the treatment effect heterogeneity across the queues to allocate resources more efficiently.

## 5.2 HMIS Data of Youth Experiencing Homelessness


We now showcase the performance of our approach to design policies that allocate resource among the U.S. homeless youth. We defer the details on data preparation to the Appendix.

## 5.3 Data Pre-Processing and Estimation

**Outcome Definition.** We focus on the likelihood of stable exit from homelessness. An exit from the system can be to any of the follow-



**(a) Success definition flow chart.**



**(b) Heterogeneous treatment effect using DR method.**

**Figure 4: HMIS data.**

ing destinations: “family,” “self-resolved,” “RRH,” “PSH,” “deceased,” or “incarcerated.” Exiting due to incarceration or being deceased are undesirable outcomes and are encoded as  $Y = 0$  (left branch). “Family,” “self-resolve,” “RRH,” and “PSH” are desirable outcomes but may be temporary exits, meaning that the individual may return to homelessness shortly after. In addition, there are recorded exits that are simply due to a “move” in the system from one service to another. We distinguish between these cases by checking whether an individual is “still housed,” i.e., is at the stable exit destination. If re-housed, we consider a 30-day threshold to decide whether it is a return to homelessness ( $Y = 0$ ) or a move in the system ( $Y = 1$ ). This procedure for defining outcome is summarized in Figure 4a.

**Propensity Estimation.** In order to obtain an unbiased estimation of the policy value, IPW and DR approaches rely on propensity values. In our setting, the propensities are unknown but can be estimated from data. This poses a challenge to find a model that fits the data while being well-calibrated. We use different statistical models for multi-class classification to estimate  $\mathbb{P}(R = r|X = \mathbf{x})$ ,  $\forall r \in \mathcal{R}$ . We evaluate models based on the predictive power, calibration, and fairness. For fairness, we adopt the test fairness criteria in [12] since evaluating the policy value across different protected groups requires propensity values that are well-calibrated for those groups. We defer the details on model selection to the Appendix. We note that the original dataset does not satisfy the positivity assumption. That is, some groups of individuals have only received a subset of the resources. Therefore, for data points with propensities less than 0.001, we follow the status-quo policy and we exclude them from the policy optimization.



Policy	Rates of Stable Exit from Homelessness			Wait Time (days)
	CT	DM	DR	
OPT	0.76	0.74	0.75	142.67
OPT-Fair (race)	0.76	0.75	0.76	142.64
OPT-Fair (age)	0.76	0.75	0.75	142.64
FCFS	0.68	0.68	0.66	142.64
SQ	0.66	0.63	0.63	182.21
SQ (data)	0.73	0.73	0.73	156.77

**Table 1: Out-of-sample estimated policy value measured in terms of rates of stable exit from homelessness and wait times.**

*Outcome Estimation.* DM and DR methods rely on a model of the outcome under different resources. We compare an array of models in terms of accuracy, calibration, and test-fairness. The results are summarized in the Appendix.


*Heterogenous Treatment Effect Estimation.* We use causal trees with minimum node size equal to 15 to estimate the average treatment effects across the NST score range for RRH and PSH. According to Figure 4b, PSH consistently has a higher treatment effect than RRH indicating that it is a more effective resource. Further, the treatment effect of both resources increase with score which suggests that higher-scored individuals benefit more from these resources. We also provide results on the (unbiased) probability of exiting homelessness versus NST score in the Appendix.

*Arrival Rate.* Once the queues are constructed, we estimate the arrival rate of individuals from data. Given the heavy-traffic condition, we calculate the required rate of SO as  $\mu_{SO} = \max(\lambda_Q - \mu_{RRH} - \mu_{PSH}, 0)$ . Further, in the HMIS data the resource arrival rates vary with time. In particular, between 2016 and 2017 there is a sharp decrease in the rate of PSH and RRH. Since the rate of resources is often known a-priori to the organisations, in the test data we re-evaluate the arrival rates and re-optimize according to those parameters.

## 5.4 Policy Optimization Results


We now present the policy optimization results along three distinct objectives: policy value measured in terms of rate of stable exit from homelessness, fairness by race and age, and wait time. Table 1 summarizes the results, where OPT is the optimal policy value without fairness constraints and OPT-fair (race), and OPT-fair (age) represent our method with fairness constraints over race and age, respectively. As baselines, we simulate both a fully FCFS policy and the status quo policy SQ (see Figure 1). We also compare with the deployed policy in the data SQ (data). As IPW suffers in small-propensity settings, we exclude it from the estimators.

From Table 1, OPT, OPT-fair (race), and OPT-fair (age) all outperform the baseline policies. Specifically, OPT significantly improves the rate of stable exit from homelessness by 19% and 13% (under DR estimates) over SQ and FCFS policies, respectively. Perhaps surprisingly, SQ performs worse than FCFS which is due to how the cut scores are designed. According to SQ individuals with scores 4-7 are matched to RRH. However, the RRH treatment effect is highest for scores above 7 (See Figure 4b). Compared to SQ (data), our policy values are competitive. We improve the wait time over SQ and SQ (data) by 21% and 9%, respectively and obtain values similar to FCFS policy. This is because we have imposed constraints that ensure a single CRP component and subsequently minimum wait time. As




**Figure 5: Out-of-sample rates of exit from homelessness by race (left panel) and age (right panel) using the DR estimator.**

a result, further algorithmic improvement is not possible unless problem inputs, such as resource arrival rates, change. We note that it is possible to relax the wait time constraints such that the average wait time is not minimum [2]. In our setting, this did not lead to an improvement in the policy value. Finally, we observe that OPT and OPT-fair have similar policy values, indicating a small cost for fairness. This can be explained by the fact that OPT-fair optimizes over queues split by both the protected attribute and risk score, which provides more flexibility to target the resources to different protected groups despite the problem being more constrained. Figure 5 compares the worst-case rate of exiting homelessness across age (below and over 17 years old) and racial groups (White, Black,




**Figure 6: Optimal Topology**



**Figure 7: Optimal and fair matching topology by race. Individuals are divided into four different score groups:  $S < 6$ ,  $S = \{6\}$ ,  $S = \{7, 8\}$ ,  $S > 8$ . Queues are constructed based on score groups and race jointly. Solid lines indicate that a resource is connected to the entire score group (a collection of queues). Dotted lines indicate connection to a single queue within the score group. For example, SO is only connected to the individuals with  $S = \{6\}$  and race White.**

and Other) according to DR estimator and in the test data. First, we observe that an FCFS policy does not necessarily result in policies that are fair in terms of their outcomes neither by age nor by race. This is because FCFS policies ignore treatment effect heterogeneity. In other words, according to the FCFS discipline, everyone has the same probability of receiving any one of the resource types (fairness in allocation). However, not everyone benefits equally from the resources. Indeed, Black individuals seem to suffer the most under a fully FCFS policy. SQ also yields a low worst-case performance mainly due its low overall performance. SQ (data) has relatively better worst-case performance. However, there is still a significant gap between the performance of Black/Other groups and Whites. By explicitly imposing fairness constraints on policy outcomes across protected groups, OPT-Fair significantly improves the performance for the Black and Other groups. Figure 5, similar observations can be made for fairness by age, where compared to baselines, OPT-Fair exhibits significant improvements in the policy value for those with age below 17.

We now present a schematic diagram of OPT and OPT-fair matching topologies. Figure 6 is the matching topology corresponding to OPT policy. Compared to SQ, OPT uses different cut points on NST score, specifically for the lower-scoring individuals. Across the four score groups, we observe a gradual transition from eligibility for a more resource-intensive intervention (PSH) to a basic intervention (SO). Figure 7 depicts OPT-fair topology for fairness on race, in which queues are constructed using the joint values of



**Figure 8: Fair topology when queues are not divided by race.**

NST score and race. According to this figure, PSH is matched to all individuals with scores above 9 as well as mid-scoring Black individuals, i.e.,  $6 \leq \text{score} \leq 9$ . RRH is connected to every individual in the mid-score range. Our modeling strategy uses the protected characteristics in order to ensure fairness. This is motivated by discussions with our community advisory board, including housing providers/matchers and people with past history of homelessness, who suggested that in order to create a fair housing allocation system there ought to be special accommodations for historically disadvantaged people. Our policies align with affirmative action policies that take individuals' protected attributes into account in order to overcome present disparities of past practices, policies, or barriers by prioritizing resources for underserved or vulnerable groups. In this regard, recently HUD restored Affirmatively Furthering Fair Housing rule that requires "HUD to administer its programs and activities relating to housing and urban development in a manner that affirmatively furthers the purposes of the Fair Housing Act", extending the existing non-discrimination mandates [13].

Our approach can also be extended to non-affirmative policies. This is possible by imposing constraints that ensure a topology has the same connections to all protected groups within a score group. Such constraints are expressible as linear constraints and can be easily incorporated in Problem (2). We demonstrate the result for fairness on race in Figure 8. We observe that all individuals who belong to a certain queue, regardless of their race, are eligible for the same types of resources. However, as a result of combining the queues, the worst-case policy value across the racial groups decreases from 0.76 to 0.73 which still outperforms SQ and SQ (data) with worse-case value of 0.61 and 0.69, respectively. We defer the results for fairness by age to Appendix.

## 6 RESEARCH ETHICS AND SOCIAL IMPACT

Recently, there has been a significant growth in algorithms that assist decision-making across various domains [19, 30, 41, 43]. Homelessness is a pressing societal problem with complex fairness considerations which can benefit greatly from data-driven solutions. As empirical evidence on ethical side effects of algorithmic decision-making is growing, care needs to be taken to minimize the possibility of indirect or unintentional harms of such systems. We take steps towards this goal. Specifically, we propose *interpretable* data-driven policies that make it easy for a decision-maker to identify and prevent potential harms. Further, we center our development around issues of *fairness* that can creep into data from different

sources such as past discriminatory practices. We provide a flexible framework to design policies that overcome such disparities while ensuring efficient allocations in terms of wait time and policy outcome.

There are also crucial consideration before applying our framework in real-world. Our approach relies on several key assumptions about the data. Specifically, the consistency assumption requires that there is only one version of PSH, RRH, and SO. In practice, different organizations may implement different variants of these interventions. For example, combining substance abuse intervention with PSH and RRH. Such granular information about the interventions, however, is not currently recorded in the data which may impact CATE estimates. Further, the exchangeability assumption requires that there are no unobserved confounders between treatment assignment and outcomes. Even though our dataset consists of a rich set of features for each individual, in practice, unobserved factors may influence the allocation of resources which calls for more rigorous inspection of service assignment processes. Unobserved confounders may lead to biased estimates of treatment effects which in turn impacts the allocation policies. In addition, our dataset consists of samples from 16 communities across the U.S., which may not be representative of new communities or populations. Hence, the external validity of such policies should be carefully studied before applying to new populations. Finally, there are other domain-specific constraints that we have not considered as they require collecting additional data. For example, resources can not be moved between different CoCs. We leave such considerations to future work.

## ACKNOWLEDGMENTS

P. Vayanos and E. Rice gratefully acknowledge support from the Hilton C. Foundation, the Homeless Policy Research Institute, and the Home for Good foundation under the “C.E.S. Triage Tool Research & Refinement” grant. P. Vayanos is also grateful for the support of the National Science Foundation under CAREER award number 2046230. Finally, we thank the project’s community advisory board for helpful discussions throughout the development of this work.

## REFERENCES

- [1] Ivo Adan and Gideon Weiss. 2014. A skill based parallel service system under FCFS-ALIS — steady state, overloads, and abandonments. *Stochastic Systems* 4, 1 (2014), 250–299. <https://doi.org/10.1214/13-ssy117>
- [2] Philipp Afeche, René Caldentey, and Varun Gupta. 2021. On the Optimal Design of a Bipartite Matching Queueing System. *Operations Research* (2021). <https://doi.org/10.1287/opre.2020.2027>
- [3] Nick Arnosti and Peng Shi. 2019. How (Not) to Allocate Affordable Housing. *AEA Papers and Proceedings* 109 (2019). <https://doi.org/10.1257/pandp.20191031>
- [4] Nick Arnosti and Peng Shi. 2020. Design of Lotteries and Wait-Lists for Affordable Housing Allocation. *Management Science* 66, 6 (2020). <https://doi.org/10.1287/mnsc.2019.3311>
- [5] Barış Ata and Mustafa H. Tongaralak. 2013. On scheduling a multiclass queue with abandonments under general delay costs. *Queueing Systems* 74, 1 (2013), 65–104. <https://doi.org/10.1007/s11134-012-9326-6>
- [6] Mohammad Javad Azizi, Phebe Vayanos, Bryan Wilder, Eric Rice, and Milind Tambe. 2018. Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 10848 LNCS. Springer, Cham, Delft, The Netherlands, 35–51. [https://doi.org/10.1007/978-3-319-93031-2\\_3](https://doi.org/10.1007/978-3-319-93031-2_3)
- [7] Chaithanya Bandi, Nikolaos Trichakis, and Phebe Vayanos. 2019. Robust multi-class queueing theory for wait time estimation in resource allocation systems. *Management Science* 65, 1 (2019), 152–187. <https://doi.org/10.1287/mnsc.2017.2948>
- [8] Dimitris Bertsimas, Jack Dunn, and Nishanth Mundru. 2019. Optimal Prescriptive Trees. *INFORMS Journal on Optimization* 1, 2 (2019), 164–183. <https://doi.org/10.1287/ijoo.2018.0005>
- [9] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. 2013. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61, 1 (2013), 73–87. <https://doi.org/10.1287/opre.1120.1138>
- [10] Francisco Castro, Hamid Nazerzadeh, and Chiwei Yan. 2020. Matching queues with reneging: a product form solution. *Queueing Systems* 96, 3–4 (2020), 359–385. <https://doi.org/10.1007/s11134-020-09662-y>
- [11] Hau Chan, Eric Rice, Phebe Vayanos, Milind Tambe, and Matthew Morton. 2017. Evidence from the past: AI decision AIDS to improve housing systems for homeless youth. In *AAAI Fall Symposium - Technical Report*, Vol. FS-17-01 - FS-17-05. AAAI Press, Stanford University, United States.
- [12] Alexandra Chouldechova. 2017. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data* 5, 2 (2017), 153–163. <https://doi.org/10.1089/big.2016.0047>
- [13] Department of Housing and Urban Development. 2021. *Restoring Affirmatively Furthering Fair Housing Definitions and Certifications*. Technical Report. Office of Fair Housing and Equal Opportunity, HUD.
- [14] John P. Dickerson and Tuomas Sandholm. 2015. FutureMatch: Combining human value judgments and machine learning to match in dynamic environments. In *Proceedings of the National Conference on Artificial Intelligence*, Vol. 1. AAAI press, Austin, Texas, United States, 622–628.
- [15] Yichuan Ding, S. Thomas McCormick, and Mahesh Nagarajan. 2021. A fluid model for one-sided bipartite matching queues with match-dependent rewards. *Operations Research* 69, 4 (2021). <https://doi.org/10.1287/opre.2020.2015>
- [16] Miroslav Dudík, John Langford, and Hong Li. 2011. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*. Omnipress 2600 Anderson St Madison WI United States, Bellevue Washington USA, 1097–1104.
- [17] Mohammad M. Fazel-Zarandi and Edward H. Kaplan. 2018. Approximating the first-come, first-served stochastic matching model with Ohm’s law. *Operations Research* 66, 5 (2018), 1423–1432. <https://doi.org/10.1287/opre.2018.1737>
- [18] Vincent A. Fusaro, Helen G. Levy, and H. Luke Shaefer. 2018. Racial and Ethnic Disparities in the Lifetime Prevalence of Homelessness in the United States. *Demography* 55, 6 (2018), 2119–2128. <https://doi.org/10.1007/s13524-018-0717-0>
- [19] Steven N. Goodman, Sharad Goel, and Mark R. Cullen. 2018. Machine learning, health disparities, and causal reasoning. <https://doi.org/10.7326/M18-3297>
- [20] Meghan Henry, Tanya de Sousa, Caroline Roddey, Swati Gayen, Thomas Joe Bednar, and Abt Associates. 2020. *AHAR: Part 1—PIT Estimates of Homelessness in the US HUD Exchange*. Technical Report. The U.S. Department of Housing and Urban Development, Office of Community Planning and Development.
- [21] Miguel A. Hernán and James M. Robins. 2013. *Causal Inference Book*. <http://www.Hsph.Harvard.Edu/Miguel-Hernan/Causal-Inference-Book/> (2013).
- [22] C. Hill, H. Hsu, M. Holguin, M. Morton, H. Winetrobe, and E. Rice. 2021. An examination of housing interventions among youth experiencing homelessness: an investigation into racial/ethnic and sexual minority status. *Journal of Public Health* (2021). <https://doi.org/10.1093/pubmed/dfab295>
- [23] Nathanael Jo, Sina Aghaei, Andres Gomez, and Phebe Vayanos. 2021. Learning Optimal Prescriptive Trees from Observational Data. (2021).
- [24] Edward Harris Kaplan. 1984. *Managing the Demand for Public Housing*. Ph.D. Dissertation. MIT.
- [25] Moniba Keymanesh, Tanya Berger-Wolf, Micha Elsner, and Srinivasan Parthasarathy. 2021. Fairness-aware Summarization for Justified Decision-Making.
- [26] Shakeer Khan and Elie Tamer. 2010. Irregular Identification, Support Conditions, and Inverse Weight Estimation. *Econometrica* 78, 6 (2010), 2021–2042. <https://doi.org/10.3982/ecta7372>
- [27] Amanda Kube, Sanmay Das, and Patrick J. Fowler. 2019. Allocating interventions based on predicted outcomes: A case study on homelessness services. In *33rd AAAI Conference on Artificial Intelligence, AAAI 2019, 31st Innovative Applications of Artificial Intelligence Conference, IAAI 2019 and the 9th AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019*. AAAI Press, Honolulu, Hawaii, United States, 622–629. <https://doi.org/10.1609/aaai.v33i01.3301622>
- [28] Avishai Mandelbaum and Alexander L. Stolyar. 2004. Scheduling flexible servers with convex delay costs: Heavy-traffic optimality of the generalized  $c\mu$ -rule. *Operations Research* 52, 6 (2004), 836–855. <https://doi.org/10.1287/opre.1040.0152>
- [29] Norweeta Milburn, Earl Edwards, Dean Obermark, and Janey Rountree. 2021. *Inequity in the Permanent Supportive Housing System in Los Angeles: Scale, Scope and Reasons for Black Residents’ Returns to Homelessness*. Technical Report. California Policy Lab.
- [30] John Monahan and Jennifer L. Skeem. 2016. Risk Assessment in Criminal Sentencing. *Annual Review of Clinical Psychology* 12 (2016). <https://doi.org/10.1146/annurev-clinpsy-021815-092945>

- [31] Matthew H. Morton, Amy Dworsky, Jennifer L. Matjasko, Susanna R. Curry, David Schlueter, Raúl Chávez, and Anne F. Farrell. 2018. Prevalence and Correlates of Youth Homelessness in the United States. *Journal of Adolescent Health* 62, 1 (2018). <https://doi.org/10.1016/j.jadohealth.2017.10.006>
- [32] Quan Nguyen, Sanmay Das, and Roman Garnett. 2021. Scarce Societal Resource Allocation and the Price of (Local) Justice. In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Press, Virtual Conference, 5628–5636.
- [33] U.S. Dept. of Housing, Office of Policy Development Urban Development, and Research. 2007. *The applicability of housing first models to homeless persons with serious mental illness: Final report*. Technical Report. OFFICE OF POLICY DEVELOPMENT AND RESEARCH (PD&R).
- [34] Orgcode. 2015. *Transition Age Youth – Vulnerability Index – Service Prioritization Decision Assistance Tool (TAY-VI-SPDAT): Next Step Tool for Homeless Youth*. Technical Report. <http://ctagroup.org/wp-content/uploads/2015/10/Y-SPDAT-v1.0-Youth-Print.pdf>.
- [35] OrgCode. 2020. The Time Seems Right: Let’s Begin the End of the VI-SPDAT.
- [36] Aida Rahmattalabi, Shahin Jabbari, Himabindu Lakkaraju, Phebe Vayanos, Max Izenberg, Ryan Brown, Eric Rice, and Milind Tambe. 2021. Fair Influence Maximization: A Welfare Optimization Approach. In *Proceedings of 35th AAAI Conference on Artificial Intelligence*. AAAI Press, New York, United States, 11630–11638.
- [37] Aida Rahmattalabi, Phebe Vayanos, Anthony Fulginiti, Eric Rice, Bryan Wilder, Amulya Yadav, and Milind Tambe. 2019. Exploring algorithmic fairness in robust graph covering problems. In *Advances in Neural Information Processing Systems*, Vol. 32. MIT Press, Vancouver, Canada, 15776–15787.
- [38] John Rawls. 1999. *Theory of justice. Revised Edition*.
- [39] Eric Rice. 2017. Assessment Tools for Prioritizing Housing Resources for Homeless Youth.
- [40] Eric Rice, Monique Holguin, Hsun-Ta Hsu, Matthew Morton, Phebe Vayanos, Milind Tambe, and Hau Chan. 2018. Linking Homelessness Vulnerability Assessments to Housing Placements and Outcomes for Youth. *CITYSCAPE* 20, 3 (2018), 69–86.
- [41] Lisa Rice and Deidre Swesnik. 2012. Discriminatory effects of credit scoring on communities of color.
- [42] Donald B Rubin. 2005. Causal Inference Using Potential Outcomes. *J. Amer. Statist. Assoc.* 100, 469 (2005), 322–331. <https://doi.org/10.1198/016214504000001880>
- [43] Tom Simonite. 2020. Meet the secret algorithm that’s keeping students out of college.
- [44] United States Interagency Council on Homelessness. 2015. *Opening doors: Federal strategic plan to prevent and end homelessness*. Technical Report. US Interagency Council on Homelessness.
- [45] Stefan Wager and Susan Athey. 2018. Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *J. Amer. Statist. Assoc.* 113, 523 (2018), 1228–1242. <https://doi.org/10.1080/01621459.2017.1319839>

## A APPENDIX

### A.1 Supplemental Material: Proof of Proposition 2

**Proof.** We let the  $X_q = 1$  be the event where  $P(X) = q$  ( $X_q = 0$  otherwise). Using this notation, we can write:

$$\begin{aligned}
 V(\pi_M) &= \mathbb{E} \left[ \sum_{r \in \mathcal{R}} \pi(r | X) Y(r) \right] \\
 &= \sum_{q \in \mathcal{Q}} \mathbb{P}(X_q = 1) \mathbb{E} \left[ \sum_{r \in \mathcal{R}} \pi(r | X) Y(r) \middle| X_q = 1 \right] \\
 &= \sum_{q \in \mathcal{Q}} \mathbb{P}(X_q = 1) \mathbb{E} \left[ \sum_{r \in \mathcal{R}} \pi(r | X) (Y(r) - Y(0)) \middle| X_q = 1 \right] \\
 &\quad + \mathbb{P}(X_q = 1) \mathbb{E} \left[ \sum_{r \in \mathcal{R}} \pi(r | X) Y(0) \middle| X_q = 1 \right] \\
 &= \sum_{q \in \mathcal{Q}} \mathbb{P}(X_q = 1) \mathbb{E} \left[ \sum_{r \in \mathcal{R}} \pi(r | X) (Y(r) - Y(0)) \middle| X_q = 1 \right] \\
 &\quad + \mathbb{P}(X_q = 1) \mathbb{E} [Y(0) | X_q = 1] \\
 &= \sum_{q \in \mathcal{Q}} \mathbb{P}(X_q = 1) \mathbb{E} \left[ \sum_{r \in \mathcal{R}} \pi(r | X) (Y(r) - Y(0)) \middle| X_q = 1 \right] + C \\
 &= \sum_{q \in \mathcal{Q}} \mathbb{P}(X_q = 1) \sum_{r \in \mathcal{R}} \pi(r | X) \mathbb{E} [Y(r) - Y(0) | X_q = 1] + C \\
 &= \sum_{q \in \mathcal{Q}} \frac{\lambda_q}{\lambda_Q} \sum_{r \in \mathcal{R}} \frac{f_{qr}}{\lambda_q} \tau_{qr} + C \\
 &= \sum_{q \in \mathcal{Q}} \sum_{r \in \mathcal{R}} \frac{f_{qr} \tau_{qr}}{\lambda_Q} + C,
 \end{aligned}$$

where  $C = \sum_{q \in \mathcal{Q}} \mathbb{P}(X_q = 1) \mathbb{E} [Y(0) | X_q = 1] = \mathbb{E} [Y(0)]$ . ■

### A.2 Supplemental Material: Proof of Proposition 3

**Proof.** We first prove part one and show the conditional independence for each component  $Y_r$  of the potential outcome vector. The proof is in the same vein as the balancing scores in the causal inference literature which is essentially a low-dimensional summary of the feature space that facilitates causal inference for observational data in settings with many features. For binary potential outcomes, we have

$$\begin{aligned}
 \mathbb{P}(Y_r = 1 | S, R) &= \mathbb{E}[Y_r | S, R] \\
 &= \mathbb{E}[\mathbb{E}[Y_r | S, R, X] | S, R] \\
 &= \mathbb{E}[\mathbb{E}[Y_r | S, X] | S, R] \\
 &= \mathbb{E}[\mathbb{E}[Y_r | X] | S, R] \\
 &= \mathbb{E}[S_r | S, R] \\
 &= S_r,
 \end{aligned}$$

where the third line follows the assumption of the proposition and the fourth line holds since  $S$  is essentially a function of  $X$  and can

be dropped. We also show

$$\begin{aligned}
 \mathbb{P}(Y_r = 1 | S) &= \mathbb{E}[Y_r | S] \\
 &= \mathbb{E}[\mathbb{E}[Y_r | S, X] | S] \\
 &= \mathbb{E}[\mathbb{E}[Y_r | X] | S] \\
 &= \mathbb{E}[S_r | S] \\
 &= S_r.
 \end{aligned}$$

We proved  $\mathbb{P}(Y_r = 1 | S, R) = \mathbb{P}(Y_r = 1 | S)$ . We now prove the second part of the proposition.

$$\mathbb{P}(\mathbb{P}(R = r | X = \mathbf{x}) > 0) = 1 \Rightarrow \mathbb{P}(\mathbb{P}(R = r, X = \mathbf{x}) > 0) = 1$$

$$\begin{aligned}
 \mathbb{P}(\mathbb{P}(R = r, X = \mathbf{x}) > 0) &= \mathbb{P}(\mathbb{P}(R = r, X = \mathbf{x}, S = \mathbf{s}) > 0) \\
 &\leq \mathbb{P}(\mathbb{P}(R = r, S = \mathbf{s}) > 0).
 \end{aligned}$$

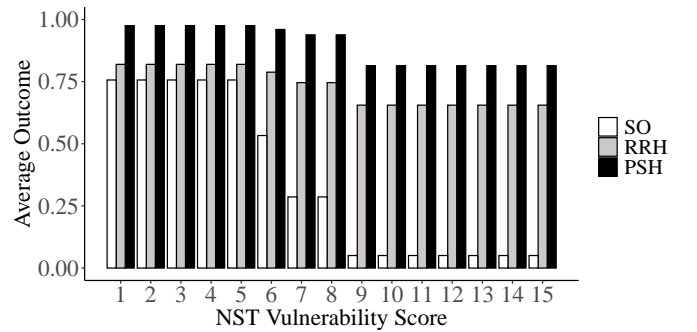
It follows that  $\mathbb{P}(\mathbb{P}(R = r, S = \mathbf{s}) > 0) = 1$  for all values of  $\mathbf{s}$ . ■

## B SUPPLEMENTAL MATERIAL: COMPUTATIONAL RESULTS

**HMIS Data Preparation.** We used HMIS dataset collected between 2015 and 2017 across 16 communities in the United States. The dataset contains 10,922 homeless youth and 3464 PSH and RRH resources combined. We removed all those with veteran status (54 data points), pending and unknown outcomes (4713 data points). We grouped Hawaiian/Pacific Islander, Native American, Hispanics, Asian under ‘Other’ category as no significant statistical inference can be made on small set of observations within each individual category. Further, we removed 6 data points with no gender information. We use a median date 08/13/2015 as the cut-off date to separate train and test sets.

**Outcome Estimation.** Figure 9 depicts the average outcome across different score values  $\mathbb{E}[Y(r) | S = s] \forall r \in \mathcal{R}$ , using the DR estimate. Under SO, after  $S = 8$ , there is a significant drop in average outcome. Average outcomes under PSH and RRH also exhibit a decline with score. However, they remain highly effective even for high-scoring youth.

**Propensity Score.** In order to evaluate different policies using IPW and DR methods, we estimated the propensity scores, i.e.,



**Figure 9: Probability of exiting homelessness across the NST score range estimated using the DR method.**

	Model	In-Sample Accuracy (%)	Out-of-Sample Accuracy (%)
NST Score	Multinomial Regression	72.5	73.7
	Neural Network	76.4	76.5
	Decision Tree	76.3	76.2
	Random Forest	76.4	76.3
All Features	Multinomial Regression	75.4	73.5
	Neural Network	80.4	77.2
	Decision Tree	79.2	78.5
	Random Forest	99.7	79.3

Table 2: Prediction accuracy for propensity estimation using HMIS data.

$\pi_0(R = r \mid X = \mathbf{x})$ . Table 2 summarizes the accuracy across different models. We consider two models, one that uses only the NST score and one that uses the entire set of features in the data. We observe that, even though the policy recommendations only use NST score, including other features help improve the accuracy. In addition, the decision tree and random forest are the top-performing models. Although random forest exhibits over-fitting (in-sample accuracy = 99.6%) its out-of-sample accuracy (79.3%) outperforms other models. In addition to accuracy, the propensity models should be well-calibrated. That is, the observed probability should match the predicted probability. We plot the reliability diagrams in Figure 10, where  $y$ -axis is the observed probability in the data and the  $x$ -axis is the predicted value. The dots correspond to values

of different bins. A well-calibrated model should lie on the  $y = x$  diagonal line.

As seen in Figure 10, random forest and neural network models have relatively better calibration property. Finally, in our model selection, we take fairness considerations into account. In particular, we study the calibration of the models across different demographic groups for which fair treatment is important. Since ultimately we use the probability estimates, not the binary prediction, it is important to ensure that across different demographic groups, the models are well-calibrated. We adopted test-fairness notion [12]. We fit a model to predict the resource one receives, based on the predicted propensities and demographic features. In a well-calibrated model across demographic groups, the coefficients of the demographic

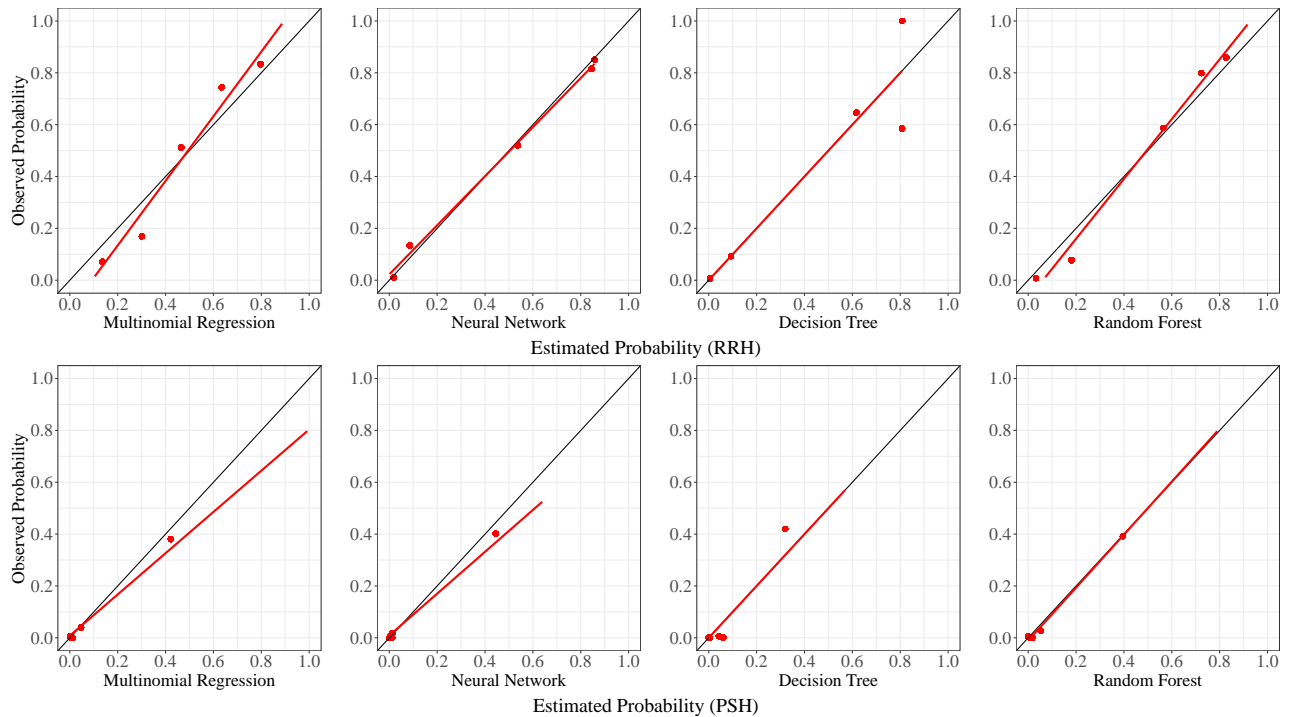


Figure 10: Reliability diagram of propensity estimation, RRH (top) and PSH (bottom).



Coeffs.	Estimates	p-value	Coeffs.	Estimates	p-value
Intercept	-0.012	0.066	Intercept	-0.054	5.7e-05
PSH pred.	<b>0.985</b>	<2e-16	RRH pred.	<b>1.125</b>	< 2e-16
Race = 2	0.011	0.204	Race = 2	-0.007	0.586
Race = 3	0.002	0.803	Race = 3	-0.014	0.394
Gender = 2	-0.006	0.469	Gender = 2	0.000	0.987
Age = 2	0.006	0.485	Age = 2	-0.003	0.813

**Table 3: Propensity calibration within group for PSH (left) and RRH (right) of random forest model. None of the coefficients of the demographic attributes are found to be significant. In addition, the coefficient associated with the predicted probability is close to 1 in both models, suggesting that the model is well-calibrated even when we control for the demographic attributes.**

attributes should not be statistically significant in the prediction. For the predicted values of the random forest model none of the demographic attributes coefficients were found to be statistically significant. In addition, the model were calibrated within groups with coefficient near 1. Regression results are summarized in Table 3. Hence, we chose random forest as the model of historical policy  $\pi_0$ .

**Outcome Estimation.** In the direct method, one estimates the (counterfactual) outcomes under different resources by fitting the regression models  $\mathbb{P}(Y | X = \mathbf{x}, R = r) \forall r \in \mathcal{R}$ . For model selection, we followed the same procedure as propensity score estimation. Table 4 summarizes the accuracy of different models for each type of resource.

Considering the reliability diagrams in Figure 11, we observe that logistic regression models are well-calibrated across different resources. We also investigated test-fairness of logistic regression where we fit the observed outcome against the predicted outcome and demographic features. Results are summarized in Table 5. As

seen, the coefficients of demographic features are not significant, suggesting that test-fairness is satisfied.

#### Optimal Matching Topology for Fairness over Age.

Figure 12 depicts the policies when fairness over age is imposed. According to this figure, across all score values youth below 17 years are eligible for PSH. On the other hand, mid- and high-scoring youth over 17 years old, are eligible for PSH. We further imposed constraints to ensure within each score group, the connections are the same for different age groups. Figure 13 illustrates the resulting matching topology, according to which individuals who score above 7 are eligible for RRH and PSH, regardless of their age. Those who score 6 are eligible for all three resource types. Finally, All youth with score below 6 are only eligible for SO. We observe that all individuals who belong to a certain queue, regardless of their age, are eligible for the same types of resources. As a result of combining the queues that depended on age, the worst-case policy value across the age groups decreased from 0.74 to 0.69 which still outperforms the SQ (data) with worst-case performance of 0.64.

	Model	PSH	RRH	SO
NST	Logistic Regression	83.1	78.8	90.0
	Neural Network	83.9	78.9	90.0
	Decision Tree	83.9	78.9	90.0
	Random Forest	83.1	78.6	90.0
NST + Demographic	Logistic Regression	83.1	78.8	90.0
	Neural Network	81.6	78.3	90.3
	Decision Tree	83.9	78.8	90.0
	Random Forest	83.9	78.1	90.0
All Features	Logistic Regression	81.9	82.2	90.3
	Neural Network	83.9	78.8	86.8
	Decision Tree	74.3	81.1	90.0
	Random Forest	83.9	81.4	90.0

**Table 4: Out-of-Sample Accuracy (%) of different outcome estimation models (outcome definition in Figure 4a).**

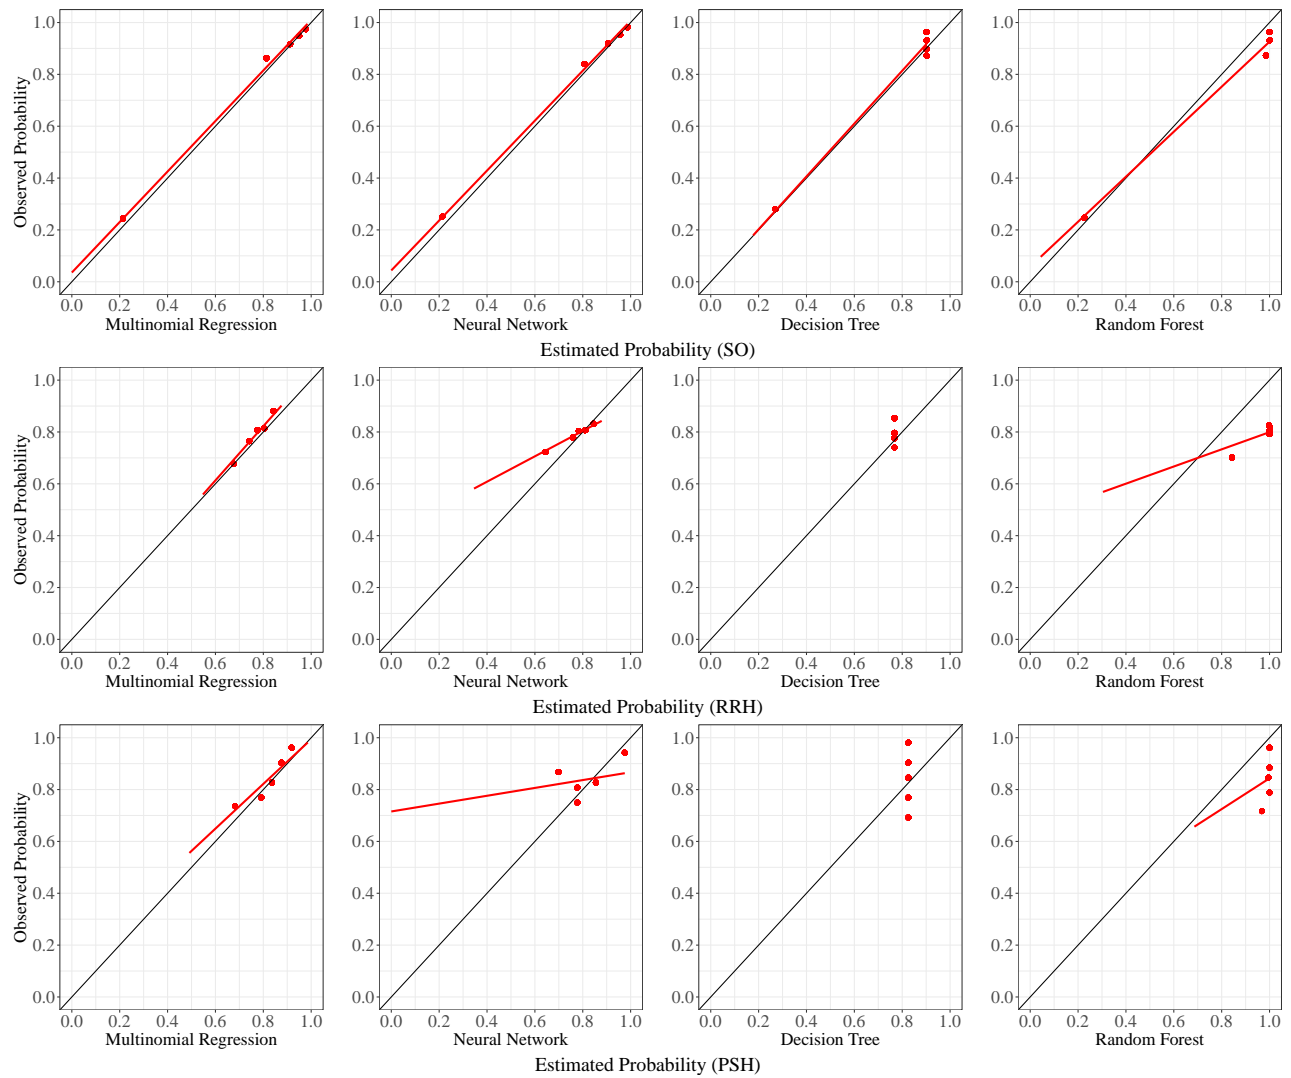




Figure 11: Reliability diagram of outcome estimation, RRH (top) and PSH (bottom).

Coeffs.	Estimates	<i>p</i> -value	Coeffs.	Estimates	<i>p</i> -value	Coeffs.	Estimates	<i>p</i> -value
Intercept	0.147	0.489	Intercept	-0.122	0.645	Intercept	0.035	0.148
PSH pred.	<b>0.853</b>	0.000	RRH pred.	<b>1.172</b>	0.000	SO pred.	<b>0.974</b>	<2e-16
Race = 2	-0.021	0.666	Race = 2	0.028	0.386	Race = 2	-0.000	0.973
Race = 3	-0.061	0.324	Race = 3	0.025	0.504	Race = 3	0.023	0.226
Gender = 2	0.003	0.954	Gender = 2	-0.021	0.433	Gender = 2	-0.008	0.618
Age = 2	0.079	0.202	Age = 2	0.003	0.931	Age = 2	-0.011	0.542

Table 5: Outcome calibration within group for PSH (left) and RRH (right) of logistic regression model. None of the coefficients of the demographic attributes are found to be significant. In addition, the coefficient associated with the predicted probability is close to 1 in both models, suggesting that the model is well-calibrated even when we control for the demographic attributes.



**Figure 12: The matching topology split by resource type: left (SO), middle (RRH) and right (PSH). The solid line indicates that the resource is connected to the entire queue. The dotted line indicates connection to a sub-group within the queue. For example, in the left figure, SO is only connected to the individuals with NST = 6 and age over 17.**



**Figure 13: Fair topology (age)**