

# Learning Optimal and Fair Policies for Online Allocation of Scarce Societal Resources from Data Collected in Deployment

Bill Tang

USC CAIS Center for Artificial Intelligence in Society, yongpeng@usc.edu

Çağıl Koçyiğit

Luxembourg Centre for Logistics and Supply Chain Management, University of Luxembourg, cagil.kocyigit@uni.lu

Eric Rice

USC CAIS Center for Artificial Intelligence in Society, ericr@usc.edu

Phebe Vayanos

USC CAIS Center for Artificial Intelligence in Society, phebe.vayanos@usc.edu

We study the problem of allocating scarce societal resources of different types (e.g., permanent housing, deceased donor kidneys for transplantation, ventilators) to heterogeneous allocatees on a waitlist (e.g., people experiencing homelessness, individuals suffering from end-stage renal disease, Covid-19 patients) based on their observed covariates. We leverage administrative data collected in deployment to design an online policy that maximizes expected outcomes while satisfying budget constraints, in the long run. Our proposed policy waitlists each individual for the resource maximizing the difference between their estimated mean treatment outcome and the estimated resource dual-price or, roughly, the opportunity cost of using the resource. Resources are then allocated as they arrive, in a first-come first-serve fashion. We demonstrate that our data-driven policy almost surely asymptotically achieves the expected outcome of the optimal out-of-sample policy under mild technical assumptions. We extend our framework to incorporate various fairness constraints. We evaluate the performance of our approach on the problem of designing policies for allocating scarce housing resources to people experiencing homelessness in Los Angeles based on data from the homeless management information system. In particular, we show that using our policies improves rates of exit from homelessness by 5.16% and that policies that are fair in either allocation or outcomes by race come at a very low price of fairness.

*Key words:* homelessness, fairness, efficiency, scarce resource allocation, data-driven optimization, causal inference

---

## 1. Introduction

We study the problem of allocating basic resources of different types to heterogeneous individuals within high-stakes social settings subject to budget and fairness constraints. We are particularly motivated by housing allocation for individuals experiencing homelessness in Los Angeles (LA)

County. According to the Los Angeles Homeless Services Authority (LAHSA), more than 75,000 people are currently unsheltered in LA County as of June 2024 (LAHSA 2024a). By comparison, the latest Housing Inventory Count in June 2024 from LAHSA reported only around 28,600 permanent housing unit interventions (LAHSA 2024b), many of which are already occupied. While permanent housing interventions are a key aspect of addressing homelessness (U.S. HUD 2007), the number of individuals experiencing homelessness far exceeds the capacity of available resources. Due to the scarcity of resources available relative to demand for them, there is a need to allocate housing in an effective way as individuals and housing resources arrive.

Additionally, since housing resources are a public good with significant impact for individuals experiencing homelessness, it is important that an allocation policy is, in some sense, fair. A lack of fairness in the system may erode the trust and effectiveness of any policy. Indeed, in our motivating example, LAHSA policymakers and community members are particularly concerned with issues of fairness in allocation and outcomes within homelessness services. A 2018 report from LAHSA’s Ad Hoc Committee on Black People Experiencing Homelessness found that while Black individuals represented 9% of the overall population in LA County, they made up 40% of the population experiencing homelessness (LAHSA 2018b). Furthermore, while housing resources are allocated at equal rates across racial groups, Black individuals receiving Permanent Supportive Housing (PSH) experienced higher rates of returns to homelessness relative to other racial groups. These findings have led to further investigations into causes of racial inequities in outcomes within the current system (Milburn et al. 2021) and a policy focus on racial equity to address existing injustices. Therefore, an allocation system needs to be not just effective, but also fair to be adopted and trusted by policymakers and community members.

Currently in the U.S., many local communities pool resources into a centralized planning system known as Continuum of Care (CoC), which coordinates housing and services funding for people experiencing homelessness. Within CoCs, Coordinated Entry Systems (CES), a network of homeless service providers and funders, use a standardized process to manage housing and supportive services and connect individuals to appropriate interventions. Individuals seeking housing in the LA CoC have various entry points into the system, such as access centers, street outreach, or emergency shelters. Following entry, they may take a self-reported survey to be assessed for eligibility and vulnerability. The assessment survey, known as Vulnerability Index–Service Prioritization Decision Assistance Tool (VI-SPDAT), consists of a series of questions to measure vulnerability, such as prior housing history and disabilities, and weights responses to output a score between 0 and 17, where a higher score indicates greater vulnerability (OrgCode 2015). Generally, the score is used in determining and prioritizing the most vulnerable individuals to receive more supportive resources. Many communities, based on OrgCode recommendations, have used score cutoffs for prioritization

or guidance in the past (LAHSA 2018a). Individuals scoring 8-17 are considered “high risk” and are prioritized for Permanent Supportive Housing (PSH), the most supportive and intensive type of housing resource. Individuals scoring 4-7 are recommended for Rapid-Rehousing (RRH), or short-term rental subsidies, while those scoring less than 4 are only eligible for services, which we refer to as Service Only (SO) (OrgCode and Community Solutions 2015). Based on these scores and the VI-SPDAT assessments, case managers, who are social service workers assisting individuals experiencing homelessness, will determine the appropriate resource to match an individual to.

While the above prioritization policy, based on just score cuts, is interpretable, both the thresholds and the weights on responses to construct the score are not tied to any data on intervention outcomes nor to arrival rates of individuals and resources. This leads to three issues. First, it is not clear that assigning the most vulnerable individuals PSH, or even RRH, would lead to the best overall outcomes, such as reducing returns to homelessness after intervention. Second, there may not be enough housing available to serve all individuals prioritized for that resource because the thresholds do not account for arrival rates. Indeed, even OrgCode, the organization that developed the VI-SPDAT, has called for alternative approaches beyond the VI-SPDAT in 2020 (OrgCode 2020). Finally, a threshold prioritization tool cannot address stakeholder concerns around racial disparities in housing outcomes. An allocation system that is equitable in outcomes would need to account for the heterogeneity of treatment effects by protected groups such as race, gender, or age (Jo et al. 2023). In contrast to the existing tool, we seek to incorporate treatment outcomes and effects into the design of a policy for maximizing desirable societal outcomes and ensuring certain notions of fairness in allocation or outcomes.

Utilizing treatment outcomes and effects in policy design, however, is challenging since they are *counterfactual*, or “what if” quantities. This means we need to know what *would happen* to an individual if they were assigned to different resources. While running randomized control trials is the gold standard for inference of causal treatment effects, in high-stakes domains such as public health and social services, experimentation may be unethical or impractical. For example, it would be unethical to deny housing from individuals experiencing dangerous and vulnerable situations. A proper experiment may also take years before conclusive results, leading to negative life consequences for many individuals during that time span. Therefore, we aim to leverage observational data for learning treatment outcomes. The historical data that LAHSA and other communities have access to was collected in deployment from administered VI-SPDAT assessments, which contains an individual’s covariates, and the LA County Homeless Management Information System (HMIS) database, which contains an individual’s trajectory of interactions with the system. These interactions include events such as engagements with street outreach workers, emergency shelter access,

and housing allocations, if any. From these trajectories, we construct a return to homelessness definition that is used as the outcome of interest.

Motivated by these issues with the existing housing allocation system, our goal is to design an efficient policy for allocating scarce housing resources to maximize positive expected outcomes, subject to budget and fairness constraints. Using historical observational data, we learn treatment outcomes based on an arriving individual’s characteristics, such as prior housing history, to appropriately match individuals to housing resources. In particular, in keeping with the existing allocation system, the assignment policy must make decisions about which resources to allocate to individuals as the resources arrive in the system, rather than waiting for a sufficient number of individuals and housing resources to aggregate in the system for offline matching. Stakeholders, including people experiencing homelessness, policy-makers, and the public at large, are dissatisfied with scarce housing sitting empty for long periods of time while thousands of people who need them sleep on the streets or in shelters each night. Furthermore, offline matching over fixed time intervals results in undesirable behavior where individuals with the same characteristics can be recommended different resources due to variations between different batches. When two individuals with the same characteristics receive different resources, this allocation does not seem fair nor interpretable to policymakers and individuals seeking assistance. This concern is captured by the notion of *individual fairness* introduced by Dwork et al. (2012), which calls for treating similar individuals similarly. A policy that assigns the same resources based on individual characteristics ensures consistency and interpretability in allocation recommendations. Additionally, implementing dynamic policies that adjust allocations based on the current state of the system is infeasible in this context, as the information necessary, such as the state of an individual in the system at any point in time, is not systematically tracked.

While we primarily focus on housing, our approach is applicable in general to allocation of scarce treatments in other public domains where system outcome efficiency and fairness are important. One example is the problem of kidney transplantations in which patients arrive to the system and join a waiting pool for a donor organ to be procured (Bertsimas et al. 2013, Dickerson and Sandholm 2015). Similar to our case, there is a significant organ shortage relative to the number of waiting patients and observational data exists from the United Network for Organ Sharing. Another example is the allocation of scarce healthcare equipment such as ventilators and critical care beds to patients arriving to the hospital with Covid-19 during spikes in hospitalizations. Observational data from electronic health records for hospitalized patients with Covid-19 can be used to learn policies for allocating equipment as patients arrive (Johnston et al. 2020, Radovanovic et al. 2021).

### 1.1. Contributions of the Paper

Our contributions in this paper are:

- We first consider policy design under budget constraints only. We propose a sample-based dual-price queuing policy by solving an offline policy learning problem using only historical (observational) data and propose an online queuing-based implementation of the derived policy. This policy establishes separate queues for each treatment, and each individual is waitlisted for the treatment that maximizes the difference between their estimated expected outcome, given their covariates, and the estimated resource dual price, which approximates the opportunity cost of assigning that treatment. Resources are allocated as they arrive in the system to the first individual in the queue for that treatment. The proposed policy ensures that capacity constraints are always met, meaning resources are assigned only when they are available in the system. We prove that, under mild technical assumptions, our policy attains the same long-run expected average outcome as the best offline policy that already knows the true distribution of covariates and (counterfactual) treatment outcomes. We establish a connection between our technical assumptions on the convergence of estimators and the standard assumptions commonly made in observational data studies, arguing that estimators satisfying our assumptions exist when the standard assumptions for observational data hold. Our framework unifies policy learning from observational data and dual-price allocation methods.
- We extend our framework to incorporate various fairness constraints, including statistical parity in allocations and outcomes between groups of a protected characteristic (e.g., race, gender, or age). We introduce fairness-constrained extensions of the sample-based dual-price queuing policy, which adjust the estimated opportunity costs for different groups using the estimated dual variables associated with the fairness constraints. While these extensions are designed to meet the respective fairness constraints in-sample, we demonstrate through numerical experiments that they can also improve fairness out-of-sample.
- We demonstrate through computational experiments on synthetic and real data that our method’s guarantees of asymptotic optimality lead to improvements in out-of-sample performance, even in finite samples. Specifically, we evaluate the policy’s performance in allocating scarce housing resources to people experiencing homelessness in Los Angeles, based on data from the HMIS database. We show that our policies improve historical rates of exits from homelessness by 5.16% (which amounts to approximately 500 additional individuals exiting homelessness per year), and policies that are fair in either allocation or outcomes by race come at a very low price of fairness.

## 1.2. Related Works

Our work is closely related to the literature on data-driven allocation of public resources and fairness, policy learning from observational data, network revenue management, and online matching.

**Data-Driven Allocation of Public Resources and Fairness.** A closely related strand of literature focuses on data-driven allocation in scarce public resource settings such as kidney allocation (Bertsimas et al. 2013, Dickerson and Sandholm 2015, Li et al. 2019), hospital ICUs (Grand-Clément et al. 2023), refugee resettlement (Bansak et al. 2018, Ahani et al. 2021, 2024, Freund et al. 2023, Bansak and Paulson 2024), or homeless services (Azizi et al. 2018, Kube et al. 2019, Rahmattalabi et al. 2022, Kaya et al. 2024, Kube et al. 2023). Our work shares similarities with several papers in this line of literature (e.g., Bansak et al. (2018), Ahani et al. (2021), Kube et al. (2023), Freund et al. (2023)) in that we use historical data to learn treatment outcomes and incorporate these estimates into resource allocation optimization problems. However, in contrast to this stream of literature, we prove the asymptotic optimality of our proposed sample-based policy, explicitly formulating the necessary properties that the estimators must satisfy and discussing their relationship to standard assumptions in observational data studies. Additionally, some of the related works, including Bansak et al. (2018) and Kube et al. (2023), formulate the corresponding resource allocation problems as mixed-integer linear optimization problems for offline matching, whereas our approach only requires solving a finite linear program to derive an online implementation. Furthermore, we consider various fairness notions, which have not been studied by the majority of these works.

The works most closely related to ours are Bertsimas et al. (2013), Li et al. (2019), Ahani et al. (2024), and Freund et al. (2023), as they, like our framework, leverage dual variables in constructing resource allocation decisions. Bertsimas et al. (2013) propose learning a scoring policy for kidney allocation that satisfies fairness constraints, but assumes knowledge of the treatment outcomes. Their scoring rule is a linear approximation of treatment outcomes minus the dual multipliers of various eligibility and fairness constraints. Li et al. (2019) and Ahani et al. (2024) propose allocations to maximize dual variable adjusted outcomes, where they estimate the dual variables based on simulated future arrival trajectories. In comparison, we only solve a single finite linear program to estimate the dual variables, provide theoretical guarantees for our proposed method, and incorporate fairness constraints. Concurrent to our work, Freund et al. (2023) tackle group fairness constraints in dynamic refugee settlement and propose dual-price assignment decisions, relying on dual variables associated with capacity and fairness constraints. They demonstrate that the regret of their proposed approach vanishes as the number of refugee cases grows. While they assume that estimates of treatment outcomes (i.e., employment probabilities) are directly available to decision-makers, our asymptotic optimality result explicitly incorporates the estimation

process from data, establishing asymptotic optimality as the number of training samples increases. Furthermore, while their fairness notion generalizes our notion of statistical parity in outcomes in Section 4.2, our notion of statistical parity in allocation in Section 4.1 addresses a distinct but complementary fairness notion. Another important difference is that we search for a policy that is a function of an individual’s covariates—therefore treating individuals with the same covariates equally—whereas they propose a matching that allows refugees with the same covariates to be treated differently.

In another related work, Bhattacharya and Dupas (2012) focus on maximizing welfare under a budget constraint for social programs with a binary treatment by learning asymptotically consistent treatment effects and propose an allocation policy based on data from a randomized experiment. In comparison to their binary treatment setting, we allow for multiple scarce treatment types, are able to impose a variety of fairness constraints, and formulate an online implementation in which an individual is matched to an appropriate resource as they arrive in the system.

Our work is also broadly related to fairness in decision-making within high-stakes domains. Many such problems incorporate fairness constraints that limit some notion of disparity between protected groups (Bertsimas et al. 2013, Corbett-Davies et al. 2017, Azizi et al. 2018, Nguyen et al. 2021, Rahmattalabi et al. 2022). Recent work has shown that different notions of fairness studied in high-stakes domains may be impossible to satisfy simultaneously (Mashiat et al. 2022, Jo et al. 2023). Other fairness concepts, e.g., based on social welfare, ordinal efficiency and envy-freeness, have been explored in algorithmic division (Rahmattalabi et al. 2021, Nguyen et al. 2021, Freeman et al. 2020, Manshadi et al. 2021) and mechanism design (Athanasoglou and Sethuraman 2011). Our work follows in the vein of imposing fairness constraints such as parity in resource allocation rates between different racial groups, as this is how policymakers within LA homeless services evaluate fairness (LAHSA 2018b).

**Policy Learning from Observational Data.** Another strand of literature focuses on policy evaluation and learning from observational data by estimating counterfactuals or correcting for selection bias in historical treatment assignments. Various works focus on constructing unbiased estimators for policy evaluation from observational data under an unconfoundedness assumption, which can then be used for policy learning. These include the direct method of using plug-in estimates of counterfactual treatment outcomes (Qian and Murphy 2011), propensity score based reweighting strategies (Swaminathan and Joachims 2015), and doubly robust combinations of outcome estimates and propensity-based reweighting (Dudík et al. 2011). Kitagawa and Tetenov (2018) and Athey and Wager (2021) study regret bounds for policy learning under propensity-weighted or doubly robust objective functions relative to some well-specified policy class with structure while Zhou et al. (2023) extend the analysis to the multiple treatment case. Other works emphasize

interpretable policy classes such as prescriptive trees (Kallus 2017, Bertsimas et al. 2019, Jo et al. 2021, Zhou et al. 2023). While these works derive important analytical results for policy classes with suitable structure, they either use a specified class like trees in practice, or suggest a general policy learning algorithm without explicitly considering how to optimize over the abstract policy class that may have budget, fairness, or functional form constraints. In contrast, we consider the general policy class of measurable functions and derive an explicit assignment policy parameterized by sample-based dual multipliers that is asymptotically optimal and can flexibly handle capacity and fairness constraints. In other words, by considering a specific policy design problem and connecting it to the dual-price controls framework, we are able to characterize an asymptotically optimal, interpretable, and computationally cheap policy across all (unrestricted) policy classes.

**Network Revenue Management.** Another related stream of literature is network revenue management (NRM), which focuses on the sale of products comprised of scarce resources to customers who request a specific product with an offer price (Talluri and Van Ryzin 2004). In domains such as airlines and hotels, the system must decide to accept or reject an arriving customer’s offer and the decision depends on the stochastic demand and revenue for products relative to resource capacities. In particular, our approach is inspired by work on bid-price controls, a class of policies where an offer is accepted only if the revenue generated is greater than the sum of optimal dual variables (bid-prices) associated with the resource capacity constraints (Talluri and Van Ryzin 1998). Early work by Talluri and Van Ryzin (1998) showed that such a control policy is asymptotically optimal as the number of requests and resources grow as long as the relative relationship of supply and demand stays constant. Since NRM problems are often formulated as dynamic programs, recent works have focused on settings where large state spaces cannot be solved tractably and computing optimal bid-prices is challenging. Tractable solutions are generally heuristics, approximate dynamic programming, or Lagrangian relaxation methods (Adelman 2007, Kunnumkal and Talluri 2016, Zhang and Weatherford 2017, Li et al. 2023). While our approach is inspired by bid-price controls within NRM, our setting differs in that the system must choose amongst a set of resources to match an arrival to rather than an accept or reject decision for a specific requested resource.

**Online Matching and Resource-Constrained Dynamic Problems.** Our work also relates to the class of online matching problems (Karp et al. 1990, Mehta et al. 2007, Aggarwal et al. 2011, Mehta et al. 2013, Devanur et al. 2013, Vera and Banerjee 2021), where one must decide which resource, among a set of capacitated resources, to immediately and irrevocably assign to each arriving individual, which results in a reward but also uses a resource. The goal is to match individuals and resources in an online fashion to maximize rewards while respecting resource capacities. Similar to our approach, online matching algorithms make matching decisions by trading off between the potential match reward and the opportunity cost of using up a scarce resource as in the case



of the AdWord problem (Mehta et al. 2007). The derivation of the tradeoff function often results from analyzing the offline or full knowledge primal-dual linear programs where the dual variables are used to parameterize the opportunity cost of using a resource. While in many online matching settings it is possible, and even desirable, for individuals with the same features to be matched to different resources, this behavior is undesirable in our setting, which is why we adopt a policy design framework rather than a matching one. When two individuals with the same features are assigned different resources, this allocation does not seem fair nor interpretable to policymakers and individuals seeking assistance. Therefore, it is important in our motivating application that two individuals with the same features receive the same type of resource.

Additionally, methodologically, our work is related to more general resource-constrained dynamic problems (Balseiro et al. 2024), which capture some of the online matching settings and network revenue management problems as special cases. Specifically, our work is closely related to the fluid certainty-equivalent control results (Balseiro et al. 2024), which leverage dual-based optimality conditions. However, our framework is based on offline policy learning translated into an online/dynamic implementation, and we work with historical observational data generated under an unknown historical policy, whereas they assume that the underlying distribution of contexts and their relation to rewards are known.

### 1.3. Organization and Notation

The remainder of this paper is structured as follows. Section 2 introduces the policy design problem and outlines the technical assumptions required for our optimality results. Section 3 studies the proposed sample-based dual-price queuing policy, along with its optimality guarantee. Section 4 studies the fairness-constrained extensions of the dual-price queuing policy. Finally, Section 5 presents numerical results based on synthetic and real data. All proofs are relegated to the online appendix.

*Notation.* All random variables in this paper are defined as measurable functions on an abstract probability space  $(\Omega, \mathcal{F}, \mathbb{P})$  and are capitalized (*e.g.*,  $X$ ) while their realizations are denoted by the same symbols in lowercase letters (*e.g.*,  $x$ ) unless stated otherwise. Throughout the paper, (in)equalities containing random variables should be interpreted to hold  $\mathbb{P}$  almost surely. We let  $\mathcal{L}_\infty(\mathcal{X}, \mathcal{C})$  denote the set of all bounded Borel-measurable functions from a Borel set  $\mathcal{X}$  to a Borel set  $\mathcal{C}$ . For any set  $\mathcal{A}$ , we use  $|\mathcal{A}|$  to denote the cardinality of the set. We use  $\mathbb{1}[\cdot]$  to denote the indicator function of a logical expression  $E$ , where  $\mathbb{1}[E] = 1$  if  $E$  is true and  $\mathbb{1}[E] = 0$  otherwise.

## 2. Problem Statement, Problem Formulation, and Assumptions

### 2.1. Problem Statement and Problem Formulation

We now formalize our problem statement. To model long-run policy performance, we consider an infinite time horizon wherein individuals indexed in the set  $\mathcal{K}$  arrive over time. Each individual

$k \in \mathcal{K}$  is characterized by their *covariates*  $X_k \in \mathcal{X} \subseteq \mathbb{R}^d$  and their *potential outcomes*  $\{Y_k^t\}_{t \in \mathcal{T}}$ , where  $\mathcal{T} = \{0, 1, \dots, m\}$  collects the treatment options, with treatment 0 representing no-treatment (or the control), and  $Y_k^t \in \mathbb{R}$  corresponds to the individual's outcome under treatment  $t$ , see e.g., Hernan and Robins (2023) for an overview of the potential outcomes framework in causal inference. For example, in the context of the housing allocation problem that motivates us, the vector of covariates  $X_k$  may collect the VI-SPDAT and HMIS information of homeless client  $k$ ; the set of treatments  $\mathcal{T}$  may collect the various housing options such as PSH, RRH, or no resource; and the potential outcome  $Y_k^t$  represents, for instance, whether client  $k$  would exit homelessness or not if assigned resource  $t$ . In particular, we regard both  $X_k$  and the  $Y_k^t$ s, observed at individual  $k$ 's *time of arrival*, as fixed quantities over time. This modeling choice is realistic in our motivating application if the time from when individuals arrive to the system and when they receive a resource is not too long. Crucially, we emphasize that an individual's covariates and potential outcomes *can change over a long period of time*, which motivates our online implementation proposed in Section 3. We numerically observe that our policy has low wait times with minimal impact on individual outcomes in Sections 5.2.4 and EC.5. Without much loss of generality we assume that  $\{(X_k, \{Y_k^t\}_{t \in \mathcal{T}})\}_{k \in \mathcal{K}}$  are independent and identically distributed (i.i.d.) across individuals and therefore may drop the index  $k$  from now on.

We do not know the joint distribution of  $(X, \{Y^t\}_{t \in \mathcal{T}})$ . However, we have access to a historical dataset of i.i.d. observations which (with a slight abuse of notation) we denote by  $\{(X_i, T_i, Y_i)\}_{i \in \mathcal{N}}$  and index in the set  $\mathcal{N}$ . Here  $X_i$  is the covariate vector of the  $i$ th observation,  $T_i$  is the historical treatment assigned to it, and  $Y_i = Y_i^{T_i}$  is the *observed outcome*, i.e., the outcome under the treatment received. Critically, we do not know and cannot control the historical treatment assignment policy and the outcomes  $Y_i^t$  for  $t \neq T_i$  remain unobserved.

We aim to use this historical observational dataset to design a policy  $\pi \in \mathcal{L}_\infty(\mathcal{X}, [0, 1]^{m+1})$  for assigning treatments to different individuals conditional on their covariates. Specifically,  $\pi^t(x)$  denotes the probability of assigning treatment  $t$  to an individual with covariates  $x$ , so that  $\sum_{t \in \mathcal{T}} \pi^t(x) = 1$  for all  $x \in \mathcal{X}$ . We focus on policies that are functions of the covariates only (being e.g., independent of the state of the system) for explainability and fairness purposes. Indeed, aligned with the notion of individual fairness (Dwork et al. 2012), such policies ensure that all individuals that are similarly situated are treated similarly and make it easier to justify why a particular allocation decision was made.

The objective of the policy design problem is to maximize the asymptotic expected average outcome of assigning treatments to arriving individuals. We can formulate the expected average outcome under a policy  $\pi$  as

$$\lim_{|\mathcal{K}| \rightarrow \infty} \mathbb{E} \left[ \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \pi^t(X_k) Y_k^t \right] = \lim_{|\mathcal{K}| \rightarrow \infty} \mathbb{E} \left[ \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \pi^t(X_k) \mathbb{E}[Y_k^t | X_k] \right] = \mathbb{E} \left[ \sum_{t \in \mathcal{T}} \pi^t(X) m^t(X) \right],$$

where  $\mathbb{E}$  denotes the expectation operator with respect to the (unknown) joint distribution  $\mathbb{P}$  of  $(X, T, \{Y^t\}_{t \in \mathcal{T}})$  and  $m^t \in \mathcal{L}_\infty(\mathcal{X}, \mathbb{R})$  is the conditional average treatment outcome under treatment  $t \in \mathcal{T}$  for covariate vector  $x \in \mathcal{X}$ , i.e.,  $m^t(x) = \mathbb{E}[Y^t | X = x]$ . The last equality follows from the definition of  $m^t$  and the assumption that individuals arrive i.i.d.

In the scarce resource settings that motivate us, each treatment type  $t \in \mathcal{T} \setminus \{0\}$  has limited availability, i.e., the arrival rate of treatments of type  $t \in \mathcal{T} \setminus \{0\}$  is less than the arrival rate of individuals. Specifically, for treatment  $t \in \mathcal{T}$ , we let  $b^t \in [0, 1]$  denote the asymptotic capacity of treatment  $t$  per individual. For example, if  $b^t = 0.3$ , then we can allocate treatment  $t$  to at most 30% of all individuals. Since treatment 0 is the no-treatment case, we set  $b^0 = 1$ . We then require policy  $\pi$  to satisfy the capacity constraints

$$\lim_{|\mathcal{K}| \rightarrow \infty} \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \pi^t(X_k) \leq b^t \quad \forall t \in \mathcal{T}.$$

As individuals arrive i.i.d. and  $\pi$  is bounded, by the strong law of large numbers we have

$$\lim_{|\mathcal{K}| \rightarrow \infty} \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \pi^t(X_k) = \mathbb{E}[\pi^t(X)] \quad \forall t \in \mathcal{T}.$$

The policy design problem is thus given by

$$\begin{aligned} z^* = \max_{\pi \in \Pi} \quad & \mathbb{E} \left[ \sum_{t \in \mathcal{T}} \pi^t(X) m^t(X) \right] \\ \text{s.t.} \quad & \mathbb{E}[\pi^t(X)] \leq b^t \quad \forall t \in \mathcal{T}, \end{aligned} \tag{P}$$

where

$$\Pi = \left\{ \pi \in L_\infty(\mathcal{X}, [0, 1]^{m+1}) : \sum_{t \in \mathcal{T}} \pi^t(x) = 1 \quad \forall x \in \mathcal{X} \right\}.$$

Without loss of generality, we take  $m^t$  to be non-negative for all  $t \in \mathcal{T}$ . In fact, if  $m^t$  did take on negative values, one could shift the values of all  $m^t$  by adding a sufficiently large constant to all of them. This does not change the optimal solution to (P) and simply shifts the optimal objective value by a constant.

Note that we cannot solve problem (P) because the functions  $m^t$  and the distribution  $\mathbb{P}$  are unknown, and we do not have access to samples from this distribution due to the observational nature of the data. Moreover, policy  $\pi$  in Problem (P) is an *offline* policy in the sense that it only depends on the personal characteristic of individuals when they arrive in the system. However, it must be implemented *online*, since in our setting, both individuals and resources arrive over time. In Section 3, we will propose a method for learning an asymptotically optimal policy for Problem (P) from the observational data. We will also introduce an online implementation of this policy that maintains the same performance. In the remainder of this section, we introduce the assumptions underlying our approach, highlighting why they are reasonable in our setting.

## 2.2. Assumptions

**2.2.1. Observational Data and Identifiability Assumptions** A fundamental challenge of policy learning from observational data is that, for any sample  $i$ , we observe the outcome  $Y_i = Y_i^{T_i}$  associated with the treatment  $T_i$  assigned by the historical policy but we do not observe the outcome that would have been obtained if any other treatment  $t \neq T_i$  had been assigned. Therefore, even evaluating the performance of a *counterfactual* policy  $\pi$  necessitates appropriate *identifiability* conditions to make causal quantities such as  $\mathbb{E}[Y^t]$  and  $m^t$  possible to estimate from the available data. We make the following assumptions that are standard within the causal inference literature (Hernan and Robins 2023).

ASSUMPTION 1. *The following conditions hold.*

- (i) *Conditional Exchangeability:*  $Y^t \perp\!\!\!\perp T \mid X \quad \forall t \in \mathcal{T}$
- (ii) *Positivity:*  $\mathbb{P}(T = t \mid X = x) > 0 \quad \forall t \in \mathcal{T} \text{ and } \forall x \in \mathcal{X}$
- (iii) *Consistency:*  $Y^T = Y$

Assumption 1 (i) says that the outcome  $Y^t$  of assigning treatment  $t \in \mathcal{T}$  and the allocated treatment  $T$  are conditionally independent given the covariates  $X = x$ . Equivalently, the joint distribution of potential outcomes is the same for all individuals in each treatment group  $t \in \mathcal{T}$  given they have the same value of  $x$ . This assumption is satisfied if we measure and condition on all covariates that are common causes of treatment assignment and outcomes. This implies that we can consider all individuals with the same  $x$  in each treatment group as exchangeable and we can calculate  $m^t(x)$  as  $\mathbb{E}[Y \mid T = t, X = x]$ . In context of our motivating example, this means the VI-SPDAT assessment responses and HMIS database contain all the relevant information for individuals experiencing homelessness that case managers use to make their final decisions on housing assignments. In our case, this assumption is likely reasonable since the VI-SPDAT is used as a prioritization tool by case managers and contains a comprehensive overview of an individual’s vulnerability factors.

Assumption 1 (ii) requires that the conditional probability of treatment assignment given features  $X = x$  is positive for each treatment. Intuitively, this means that each individual receives any particular treatment with positive probability under the historical policy. Assumption 1 (iii) states that the observed outcome under an observed treatment  $T$  is actually the potential outcome. Practically, this means that our treatments are well specified or else if there were multiple versions of a particular treatment  $t$ , then  $Y^t$  is not well defined.

**2.2.2. Conditional Average Treatment Outcomes Assumptions** Our asymptotic optimality results rely on the following technical assumption on conditional average treatment outcomes, which ensures that a dual-price queuing policy—constructed under full distributional knowledge and *not* from data—is optimal in  $(\mathcal{P})$ .

ASSUMPTION 2. *For any  $t \in \mathcal{T}$ , the following conditions hold.*

- (i) *The random variable  $m^t(X)$  is continuously distributed and has a bounded support, i.e., there exists a  $C \in \mathbb{R}_+$  such that  $|m^t(x)| \leq C$  for all  $x \in \mathcal{X}$ .*
- (ii) *For any  $t' \in \mathcal{T}$  and  $t' \neq t$ ,  $m^t(X) - m^{t'}(X)$ , the difference in conditional treatment effects between  $t$  and  $t'$ , is continuously distributed.*

Assumption 2 requires that the conditional average treatment outcomes  $m^t(X)$  for all treatments  $t \in \mathcal{T}$ , and the treatment effect differences  $m^t(X) - m^{t'}(X)$  between any two treatments  $t, t'$  are bounded and continuously distributed. Note that boundedness of  $m^t(X)$  for all  $t \in \mathcal{T}$  implies boundedness of  $m^t(X) - m^{t'}(X)$  for any  $t, t'$ . Intuitively, this assumption holds if treatment outcomes and effects take on continuous values.

The following example illustrates a simple case where the conditional average treatment outcomes are linear functions of a continuous random variable and satisfy Assumption 2.

EXAMPLE 1 (ASSUMPTIONS 2 (I) AND (II) HOLD). Suppose that there exists a single covariate  $X$  that is continuously distributed on the support  $[0, 1]$ . There are two potential treatments indexed by 0 and 1, and their respective conditional average treatment outcome functions are given by  $m^0(x) = x$  and  $m^1(x) = 2x$ . In this case, the random variables  $m^0(X)$ ,  $m^1(X)$ ,  $m^1(X) - m^0(X)$  and  $m^0(X) - m^1(X)$  are all bounded and continuously distributed.

We also show in the following two examples wherein conditions (i) and (ii) in Assumption 2 are both needed, thereby demonstrating that one does not imply the other. These also give intuition as to when Assumption 2 does not hold.

EXAMPLE 2 (ASSUMPTION 2 (I) DOES NOT IMPLY ASSUMPTION 2 (II)). Suppose again that there exists a single covariate  $X$  that is continuously distributed on the support  $[0, 1]$ . There are two potential treatments indexed by 0 and 1, and their respective conditional average treatment outcome functions are given by  $m^0(x) = x$  and  $m^1(x) = x + 1$ . The random variables  $m^0(X)$  and  $m^1(X)$  are therefore bounded and continuously distributed, which implies that Assumption 2 (i) holds. On the other hand,  $m^1(x) - m^0(x) = 1$ , and therefore  $m^1(X) - m^0(X)$  is not continuously distributed. Thus, Assumption 2 (ii) does not hold.

EXAMPLE 3 (ASSUMPTION 2 (II) DOES NOT IMPLY ASSUMPTION 2 (I)). Suppose again that there exists a single covariate  $X$  that is continuously distributed on the support  $[0, 1]$ . There are two potential treatments indexed by 0 and 1, and their respective conditional average treatment outcome functions are given by  $m^0(x) = 1$  and  $m^1(x) = x$ . Then, we have that  $m^1(x) - m^0(x) = x - 1$  and  $m^0(x) - m^1(x) = 1 - x$ , and so both  $m^1(X) - m^0(X)$  and  $m^0(X) - m^1(X)$  are bounded and continuously distributed. However,  $m^0(X)$  is constant valued in this case and not continuously distributed. Therefore, Assumption 2 (ii) does not imply Assumption 2 (i).

**2.2.3. Sample-Based Estimates Assumptions** In Section 3, we propose a sample-based dual-price queuing policy. In order to learn this policy from data, we first need to estimate the conditional treatment outcomes. Let  $\hat{m}^t(\cdot|\{(X_i, T_i, Y_i)\}_{i \in \mathcal{N}}) \in \mathcal{L}_\infty(\mathcal{X}, \mathbb{R})$  denote an estimator of the conditional mean treatment response  $m^t$  from the data. For ease of notation, we suppress the dependency of  $\hat{m}^t$  on the historical samples and simply write  $\hat{m}_n^t$  to stress that we use a sample of size  $n$ . The following assumption formally states properties we need our estimators to have and how “good” they need to be asymptotically.

ASSUMPTION 3. *For any  $t \in \mathcal{T}$ , the following condition holds.*

- (i) *The estimate  $\hat{m}_n^t$  is measurable and has bounded support, i.e., there exists a  $\hat{C} \in \mathbb{R}_+$  such that  $|\hat{m}_n^t(x)| \leq \hat{C}$  for all  $x \in \mathcal{X}$  and  $n \in \mathbb{N}$ .*
- (ii) *The estimate  $\hat{m}_n^t$  converges uniformly and almost surely to  $m^t$ , i.e.,  $\lim_{n \rightarrow \infty} \sup_{x \in \mathcal{X}} |\hat{m}_n^t(x) - m^t(x)| = 0$ .*

Assumption 3 (i) is a technical assumption needed for our theoretical results and is standard within other policy learning works as well. Assumption 3 (ii), the uniform convergence of  $\hat{m}_n^t$  to  $m^t$  for each treatment, is necessary for the asymptotic optimality of the sample-based dual-price policy; see Theorem 1. This is because we do not make any functional assumptions about the conditional expected outcomes  $m^t$ ,  $t \in \mathcal{T}$ . In fact, if one has knowledge of the form of  $m^t$ , for example that it is a linear function with parameters  $\theta$ , then Assumption 3 can be weakened to uniform convergence of the estimator parameters  $\theta$ . In our general setting, the asymptotic convergence guarantees we require in Assumption 3 (ii) are satisfied by non-parametric estimators like nearest-neighbors, kernel regression, and sieve estimators (Mack and Silverman 1982, Cheng 1984, Liero 1989, Hansen 2004, Chen and Christensen 2015, Xie 2023) if standard assumptions on observational data also hold. Specifically, under Assumption 1, the conditional average treatment outcomes  $m^t$  for  $t \in \mathcal{T}$ , which are *causal* quantities, can be expressed and estimated as *statistical* quantities that are solely a function of the observed data (Hernan and Robins 2023). When Assumption 1 holds and we can estimate a causal quantity as a statistical quantity from observed data, then the aforementioned non-parametric estimators will satisfy the uniform convergence property of Assumption 3 (ii) (Liero 1989).

### 3. Dual-Price Queuing Policy: Solution Approach and Online Implementation

In this section, we characterize an online policy, which we call the sample-based dual-price queuing policy, and show that it asymptotically attains the optimal value  $z^*$  of problem  $(\mathcal{P})$  as the number of data samples tends to infinity. To that end, we need the conditional average treatment outcome functions  $m^t$  for each  $t \in \mathcal{T}$ , which we estimate from our observational dataset. There

exist a variety of methods within the causal inference literature for learning conditional average treatment effects and outcomes such as metalearning, double machine learning, and doubly robust learning (Chernozhukov et al. 2018, Künzel et al. 2019, Foster and Syrgkanis 2023, Kennedy 2023). Since these methods allow for arbitrary estimators to be used for modeling outcomes and treatment assignments, which are components used to adjust for observational selection bias in the estimation procedure, we can use nonparametric estimators that satisfy the theoretical asymptotic assumptions we make in Assumption 3.

To showcase the approach that yields the characterization of our policy, consider the following sample approximation of problem  $(\mathcal{P})$ , where we replace the expectations with sample averages and  $m^t$ 's with their estimators  $\hat{m}_n^t$ 's.

$$\begin{aligned} \hat{z}_n^* = \max_{\pi \in \Pi} & \quad \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) \hat{m}_n^t(x_i) \\ \text{s.t.} & \quad \frac{1}{n} \sum_{i=1}^n \pi^t(x_i) \leq b^t \quad \forall t \in \mathcal{T} \end{aligned} \quad (\hat{\mathcal{P}})$$

The Lagrangian dual of this problem is given by

$$\begin{aligned} \hat{\nu}_n^* &= \min_{\mu \in \mathbb{R}_+^{m+1}} \max_{\pi \in \Pi} \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) \hat{m}_n^t(x_i) + \sum_{t \in \mathcal{T}} \mu^t \left( b^t - \frac{1}{n} \sum_{i=1}^n \pi^t(x_i) \right) \\ &= \min_{\mu \in \mathbb{R}_+^{m+1}} \max_{\pi \in \Pi} \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) (\hat{m}_n^t(x_i) - \mu^t) + \sum_{t \in \mathcal{T}} \mu^t b^t, \end{aligned} \quad (1)$$

where  $\mu$  collects the dual variables of the capacity constraints. By weak duality, we have  $\hat{\nu}_n^* \geq \hat{z}_n^*$ . For any  $\mu \in \mathbb{R}_+^{m+1}$ , the inner maximization problem is solved by the policy

$$\pi^t(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} (\hat{m}_n^{t'}(x) - \mu^{t'}) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T}, \forall x \in \mathcal{X}. \quad (2)$$

We suppress the dependence of policy  $\pi$  on  $\mu$  notationally in order to avoid clutter. We use a lexicographic tie-breaker in characterizing (2). We emphasize that our results do not rely on this particular tie-breaking rule. By substituting policy (2) into the dual problem (1), we obtain

$$\hat{\nu}_n^* = \min_{\mu \in \mathbb{R}_+^{m+1}} \frac{1}{n} \sum_{i=1}^n \max_{t \in \mathcal{T}} (\hat{m}_n^t(x_i) - \mu^t) + \sum_{t \in \mathcal{T}} \mu^t b^t. \quad (\hat{\mathcal{D}})$$

Problem  $(\hat{\mathcal{D}})$  is a convex optimization problem and can be reformulated as a finite linear program that can easily be solved by off-the-shelf solvers. Denote by  $\hat{\mu}_n^*$  an optimal solution to problem  $(\hat{\mathcal{D}})$ , which exists under Assumption 3 (i) because we can restrict the feasible set of  $(\hat{\mathcal{D}})$  to a compact set without loss of generality; see Lemma EC.4. We define a candidate policy  $\hat{\pi}_n^*$  through

$$\hat{\pi}_n^{*,t}(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} (\hat{m}_n^{t'}(x) - \hat{\mu}_n^{*,t'}) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T}, \forall x \in \mathcal{X}. \quad (3)$$

Note that policy  $\hat{\pi}_n^*$  is obtained by choosing  $\mu$  as  $\hat{\mu}_n^*$  in (2).

Next, we introduce an online implementation of policy  $\hat{\pi}_n^*$  that takes into account the fact that individuals arrive sequentially and that treatments become available over time. By construction, this policy *always* satisfies the capacity constraints, even out-of-sample and in finite horizon.

**DEFINITION 1 (SAMPLE-BASED DUAL-PRICE QUEUING POLICY).** Establish  $m$  queues each associated with a treatment  $t \in \mathcal{T} \setminus \{0\}$ . At any time, individuals in queue  $t$  are waitlisted for treatment  $t$ . There are two events triggering an action: (1) arrival of an individual, (2) arrival of a treatment.

- (1) When an individual with covariate vector  $x$  arrives, compute the treatment  $t = \min \arg \max_{t' \in \mathcal{T}} (\hat{m}_n^{t'}(x) - \hat{\mu}_n^{*,t'})$ , where  $t$  attains the maximum difference between the estimated conditional mean treatment outcome and the associated sample optimal dual variable. If  $t \neq 0$ , assign them to queue  $t$ , otherwise assign them to no-treatment. If they are the only individual waiting in queue  $t$ , check whether there is an available treatment of type  $t$ . If there is, assign treatment  $t$  to the individual.
- (2) When a treatment of type  $t$  becomes available, assign this treatment to the first individual in the queue  $t$ . If there is no one in the queue  $t$ , keep the treatment until an individual is assigned to this queue.

In finite horizon, the sample-based dual-price queuing policy assigns  $t = 0$  to individuals still waiting in queues at the end of the time horizon. Since in practice our experiments can only occur in the finite horizon setting, we will use a ‘long’ test horizon in the synthetic data experiments as a proxy for the asymptotic environment. We also note that in general, even if the historical samples are fixed,  $\hat{\mu}_n^*$ , the optimal solution to problem  $(\hat{\mathcal{D}})$ , need not be unique since the objective function of  $(\hat{\mathcal{D}})$  may not be strictly convex and therefore the sample-based dual-price queuing policy may not be uniquely specified.

In the following, we denote by  $\hat{z}_n^D$  the expected average outcome of the sample-based dual-price queuing policy, where the expectation is taken with respect to the distribution of covariates and outcomes of the individuals arriving during implementation and *not* with respect to the distribution of the historical samples. Note that  $\hat{z}_n^D$  is a random value because it relies on the historical samples, which are random. The next theorem shows that the sample-based dual-price queuing policy attains the optimal value  $z^*$  of problem  $(\mathcal{P})$  as the number of data samples tends to infinity.

**THEOREM 1.** *Under Assumptions 2 and 3, the expected average outcome  $\hat{z}_n^D$  of the sample-based dual-price queuing policy is almost surely asymptotically at least as high as the optimal value  $z^*$  of problem  $(\mathcal{P})$  as  $n \rightarrow \infty$ , that is,*

$$\limsup_{n \rightarrow \infty} z^* - \hat{z}_n^D \leq 0.$$



To prove Theorem 1, we first construct a dual-price queuing policy assuming that the joint distribution of the covariates  $X$  and the potential outcomes  $Y^t$ ,  $t \in \mathcal{T}$ , is known. We show that this dual-price queuing policy is optimal in problem  $(\mathcal{P})$ ; see EC.1 in the online appendix for further details about this construction and the corresponding result. This first part of the proof relies on establishing some structural properties of problem  $(\mathcal{P})$  and applying standard techniques based on optimality conditions from network revenue management (Talluri and Van Ryzin 1998), and more generally, from capacity-constrained dynamic programming (Balseiro et al. 2024). The idea for the rest of the proof is that if the allocations of a sample-based dual-price queuing policy asymptotically match those of a dual-price queuing policy almost surely, then the sample-based dual-price queuing policy will asymptotically generate expected average outcomes at least as good as those of a dual-price queuing policy almost surely. Since the dual-price queuing policy is optimal in problem  $(\mathcal{P})$ , then asymptotically  $\hat{z}_n^D$  is least as high as  $z^*$  almost surely. We show in the online appendix that the optimal solutions to problem  $(\hat{\mathcal{D}})$ ,  $\hat{\mu}_n^*$ , will be arbitrarily close to an optimal solution  $\mu^*$  of the dual problem of  $(\mathcal{P})$  as the number of samples increases,  $n \rightarrow \infty$ . This will imply that the allocations of a sample-based dual-price queuing policy and dual-price queuing policy match almost surely.

#### 4. Fairness Extensions

In high-stakes settings, the allocation of scarce resources needed to satisfy basic needs raises issues of fairness with respect to some protected group(s). This is especially the case when minority groups historically discriminated against due to race, gender, or other protected characteristics disproportionately make up the population supported by the allocation system. In our motivating example, fairness and equity, or the lack thereof, are key concerns for policymakers and community members (see Section 1).

In this section, we showcase fairness-constrained extensions of the sample-based dual-price queuing policy introduced in Section 3 based on two fairness notions of statistical parity in allocation and statistical parity in outcomes. In the following subsections, we introduce the aforementioned fairness notions and present the resulting fair policies. Details of the derivations of the respective policies are relegated to EC.3. The proposed extensions and their derivations can be similarly derived for other fairness notions that can be defined as linear constraints such as conditional statistical parity, which, for example, could require that the percentage of housing resources allocated to individuals with low VI-SPDAT scores should not exceed a predefined threshold. Alternatively, one could incorporate a fairness constraint that enforces lower bounds on the percentage of resources allocated to specific protected groups, which, for example, could require that the percentage of housing allocated to minority groups should be at least as high as the percentage allocated historically. For a comprehensive overview of possible fairness notions and constraints for scarce resource

allocation systems and impossibility results between these various fairness notions, see Jo et al. (2023).

Recall from Section 2 that  $X$  can include protected features  $G \in \mathcal{G}$  that represent sensitive attributes (e.g., race, gender, sex), i.e.,  $X = (X^{-G}, G)$ , where  $X^{-G}$  collects all features of  $X$  excluding  $G$ . Therefore we note that the protected feature  $G$  is also an input into our policy  $\pi$  in the following subsections.

#### 4.1. Statistical Parity in Allocation

*Statistical parity in allocation*, which requires that the treatment allocation probabilities are similar across protected groups, is an intuitive and commonly used notion by stakeholders for evaluating the fairness of high-stakes allocation systems (Jo et al. 2023). Formally, a policy  $\pi$  satisfies statistical parity in allocation if

$$\mathbb{E}[\pi^t(X) | G = g] - \mathbb{E}[\pi^t(X) | G = g'] \leq \delta \quad \forall g, g' \in \mathcal{G}, g \neq g', \forall t \in \mathcal{T}, \quad (4)$$

which requires that the allocation probabilities of any treatment  $t \in \mathcal{T}$  are approximately equal (up to a specified tolerance level  $\delta$ ) across different sensitive groups. Problem  $(\mathcal{P})$  can be extended to include the additional constraint (4). Similarly, the sample approximation  $(\hat{\mathcal{P}})$  of  $(\mathcal{P})$  can be extended to include a sample approximation of the constraint (4), i.e.,

$$\frac{1}{n_g} \sum_{i=1}^n \mathbb{1}[g_i = g] \pi^t(x_i) - \frac{1}{n_{g'}} \sum_{i=1}^n \mathbb{1}[g_i = g'] \pi^t(x_i) \leq \delta \quad \forall g, g' \in \mathcal{G}, g \neq g', \forall t \in \mathcal{T}, \quad (5)$$

where the expectations are replaced with sample averages, and  $n_g$  denotes the number of individuals with protected feature  $g$  in the data.

As in Section 3, let  $\mu \in \mathbb{R}_+^{m+1}$  collect the dual variables of the capacity constraints in  $(\hat{\mathcal{P}})$ . We now denote by  $\lambda^t(g, g') \in \mathbb{R}_+$  the dual variable of the fairness constraint (5) for the pair  $(g, g')$  and for treatment  $t$ . Following similar steps to those in Section 3, we characterize a policy

$$\hat{\pi}_{\text{alloc}, n}^{*, t}(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} \left( \hat{m}_n^{t'}(x) - \hat{\mu}_n^{*, t'} - \frac{n}{n_g} \hat{\gamma}_n^{*, t'}(g) \right) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T}, x \in \mathcal{X},$$

where for each  $g \in \mathcal{G}$  and  $t \in \mathcal{T}$ ,  $\hat{\gamma}_n^{*, t}(g) = \sum_{g' \in \mathcal{G}, g \neq g'} (\hat{\lambda}_n^{*, t}(g, g') - \hat{\lambda}_n^{*, t}(g', g))$  is an aggregation of dual variables across all  $g' \neq g$ , and  $\hat{\mu}_n^{*, t}$  and  $\hat{\lambda}_n^{*, t}(g, g')$  are optimal solutions to the Lagrangian dual problem of  $(\hat{\mathcal{P}})$  with additional fairness constraints (5). Intuitively, if  $\hat{\mu}_n^{*, t}$  is roughly interpreted as an opportunity cost/threshold for assigning treatment  $t$ , then our new policy uses  $\hat{\gamma}_n^{*, t}(g)$  to adjust that threshold for each group to achieve allocation parity. We can see that the adjustment  $\hat{\gamma}_n^{*, t}(g)$  is unrestricted in sign so that it can appropriately increase (to allocate less) or decrease (to allocate more) the ‘opportunity cost’ of each treatment for each group. Finally, our adjustment is also

inversely weighted by the population proportion of each group  $g$  to appropriately account for the protected group distribution within the population.

Using the definition of  $\hat{\pi}_{\text{alloc},n}^*$ , we can construct a fairness-constrained sample-based dual-price queuing policy similarly to the Section 3. The fairness-constrained sample-based dual-price queuing policy assigns an individual with covariates  $x$  to the queue for treatment  $t = \min \arg \max_{t' \in \mathcal{T}} \left( \hat{m}_n^{t'}(x) - \hat{\mu}_n^{*,t'} - \frac{n}{n_g} \hat{\gamma}_n^{*,t'}(g) \right)$ .

In certain allocation settings, policymakers may want to prioritize minority groups historically disadvantaged by racial, gender, or other forms of discrimination for allocation as a solution to mitigate inequities. While statistical parity in allocation ensures equal opportunities across protected groups, it may be insufficient for repairing existing disparities. Whether this prioritization is appropriate again depends on the fairness goals policymakers want to achieve. In this case, we can define a set  $\mathcal{G}^{\min}$  of minority groups and a set  $\mathcal{G}^{\text{maj}}$  of majority groups such that  $\mathcal{G}^{\min}$  and  $\mathcal{G}^{\text{maj}}$  are disjoint and  $\mathcal{G}^{\min} \cup \mathcal{G}^{\text{maj}} = \mathcal{G}$ . The goal of prioritizing allocation to minority groups could be formulated through the constraint

$$\mathbb{E}[\pi^t(X) | G = g] \leq \mathbb{E}[\pi^t(X) | G = g'] \quad \forall g \in \mathcal{G}^{\text{maj}}, g' \in \mathcal{G}^{\min}, \forall t \in \mathcal{T} \setminus \{0\}, \quad (6)$$

where the average treatment assignment probability for any minority group  $g' \in \mathcal{G}^{\min}$  is at least as high as the average treatment assignment probability for any majority group  $g \in \mathcal{G}^{\text{maj}}$ . We see that constraint (6) is a variant of (4), where we have excluded the no-treatment option of  $t = 0$ , set  $\delta = 0$ , and we enforce the constraint only for the pairs  $(g, g')$  such that  $g \in \mathcal{G}^{\text{maj}}$  and  $g' \in \mathcal{G}^{\min}$ . We can again write a sample approximation of constraint (6), include it in  $(\hat{\mathcal{P}})$ , and associate dual variables  $\lambda^t(g, g') \in \mathbb{R}_+$ , for  $g \in \mathcal{G}^{\text{maj}}, g' \in \mathcal{G}^{\min}, t \in \mathcal{T} \setminus \{0\}$ , with these constraints. Following similar steps as before, we can characterize a minority prioritizing allocation policy  $\hat{\pi}_{\text{alloc},n}^*$  as follows. For any  $x \in \mathcal{X}$  such that  $g \in \mathcal{G}^{\min}$ , we have

$$\hat{\pi}_{\text{alloc},n}^{*,t}(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} \left( \hat{m}_n^{t'}(x) - \hat{\mu}_n^{*,t'} + \frac{n}{n_g} \sum_{g' \in \mathcal{G}^{\text{maj}}} \hat{\lambda}_n^{*,t'}(g', g) \right) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T},$$

and for any  $x \in \mathcal{X}$  such that  $g \in \mathcal{G}^{\text{maj}}$ , we have

$$\hat{\pi}_{\text{alloc},n}^{*,t}(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} \left( \hat{m}_n^{t'}(x) - \hat{\mu}_n^{*,t'} - \frac{n}{n_g} \sum_{g' \in \mathcal{G}^{\min}} \hat{\lambda}_n^{*,t'}(g, g') \right) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T}.$$

For  $t = 0$ , we define  $\hat{\lambda}_n^{*,0}(g, g') = 0$  for any pair of  $g$  and  $g'$ . In this case we see that the adjustments for all minority groups  $g \in \mathcal{G}^{\min}$ , given by  $\sum_{g' \in \mathcal{G}^{\text{maj}}} \hat{\lambda}_n^{*,t'}(g', g)$ , are non-negative, and serve only to decrease the ‘threshold’ of treatment  $t$  for group  $g$ , therefore leading to more assignments. On the other-hand, since  $\sum_{g' \in \mathcal{G}^{\min}} \lambda^{*,t'}(g, g')$  is also non-negative but has opposite sign, it serves to increase the ‘threshold’ of treatment  $t$  for all  $g \in \mathcal{G}^{\text{maj}}$ , thereby leading to fewer assignments. Similarly to before, we can define a corresponding minority allocation prioritization sample-based dual-price queuing policy.

## 4.2. Statistical Parity in Outcomes

While statistical parity in allocation ensures similar chances of receiving each treatment across different protected groups, it may not be adequate for achieving fair outcomes for the individuals. If minority groups are more vulnerable to homelessness due to discrimination, then statistical parity in allocation may not address existing outcome inequities. This is precisely the case and a major concern in our motivating example, where PSH is allocated at nearly equal rates to all racial groups, but Black individuals experience higher rates of returns to homelessness compared to other racial and ethnic groups (LAHSA 2018b). For this reason, we will consider an alternative fairness notion called *statistical parity in outcomes*, which requires that a policy  $\pi$  satisfies constraints of the form

$$\mathbb{E} \left[ \sum_{t \in \mathcal{T}} \pi^t(X) m^t(X) \mid G = g \right] - \mathbb{E} \left[ \sum_{t \in \mathcal{T}} \pi^t(X) m^t(X) \mid G = g' \right] \leq \delta \quad \forall g, g' \in \mathcal{G}, g \neq g'. \quad (7)$$

Constraint (7) ensures that the expected outcomes under the allocation policy  $\pi$  are approximately equal across different protected groups. The sample approximation of constraint (7) takes the following form.

$$\frac{1}{n_g} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \mathbb{1}[g_i = g] \pi^t(x_i) \hat{m}_n^t(x_i) - \frac{1}{n_{g'}} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \mathbb{1}[g_i = g'] \pi^t(x_i) \hat{m}_n^t(x_i) \leq \delta \quad \forall g, g' \in \mathcal{G}, g \neq g' \quad (8)$$

Problem  $(\hat{\mathcal{P}})$  can be extended to include this constraint. As before, we associate dual variables with constraints of this extended problem and denote by  $\lambda(g, g') \in \mathbb{R}_+$  the dual variable of the statistical parity in outcome constraint (8) for the group pair  $g$  and  $g'$ . We denote by  $\hat{\mu}_n^*$  and  $\hat{\lambda}_n^*$  the optimal solutions to the Langragian dual problem of  $(\hat{\mathcal{P}})$  with additional statistical parity in outcomes constraints (8). Then, we characterize a modified policy

$$\hat{\pi}_{\text{out},n}^{*,t}(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} \left( \hat{m}_n^{t'}(x) \left( 1 - \frac{n}{n_g} \hat{\gamma}_n^*(g) \right) - \hat{\mu}_n^{*,t'} \right) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T}, \forall x \in \mathcal{X}, \quad (9)$$

where  $\hat{\gamma}_n^*(g) = \sum_{g' \in \mathcal{G}, g' \neq g} \left( \hat{\lambda}_n^*(g, g') - \hat{\lambda}_n^*(g', g) \right)$ . Compared to statistical parity in allocation, we see now that the adjustment is directly applied to the contribution value of the estimated conditional mean treatment outcome function  $\hat{m}_n^t$ . The adjustment  $\hat{\gamma}_n^*(g)$ , which is unrestricted in sign and inversely weighted by the population proportion of group  $g$ , will scale up or down  $\hat{m}_n^t(x)$  for a given treatment  $t$  in order for assignments to achieve parity in outcomes. If on average group  $g$ 's outcomes are not as high as other groups' outcomes, we can upweight group  $g$ 's conditional mean treatment outcomes so that group  $g$  individuals are more likely to 'pass' the 'threshold' value  $\hat{\mu}_n^{*,t}$ . This leads to individuals in group  $g$  being more likely to receive treatments under the statistical parity in outcomes policy, and therefore higher group-wise outcomes. As in the case of statistical parity in allocation in Section 4.1, we can similarly define a minority prioritization version of (9) and the respective fairness-constrained sample-based dual-price queuing policies.

## 5. Empirical Results

In this section, we evaluate the empirical performance of our methodology in two sets of experiments: (1) a synthetic example, where the data generation models for covariates, treatment assignments, and potential outcomes are known; (2) an example based on real HMIS and VI-SPDAT data from our central motivating application of allocating scarce housing resources to people experiencing homelessness. In both cases, we split the available data into training and testing/test sets, where the training data is used to learn the dual prices and counterfactual outcomes and the testing data is used to evaluate policy performance.

### 5.1. Synthetic Data Experiments

We use a simple synthetic example to evaluate the performance of our approach in a setting where counterfactuals are known and to study the effects of bias from insufficient training data, large noise terms, and model misspecification on the out-of-sample performance of our methodology. For the synthetic experiments, our primary performance metric is the ratio of the test set outcomes of the sample-based dual-price queuing policy to the outcomes of the perfect foresight test set policy (that knows both the people and resources arriving and the counterfactual outcomes). Thus, a ratio of 1 corresponds to the best possible performance. We also study the convergence of the performance of our policy to that of the perfect foresight test set policy as the training sample size is increased.

**5.1.1. Data Generation.** For the synthetic data experiments, we adapt the data generation process of Jo et al. (2021) by increasing the number of treatments from two to three to test our methodology in the multiple treatments case. In this setting, there are three possible treatments, i.e.,  $\mathcal{T} = \{0, 1, 2\}$ , and each individual is characterized by two covariates, each following a standard normal distribution, i.e.,  $X = [X^1, X^2]$  and  $X^1, X^2 \sim \mathcal{N}(0, 1)$ . The potential outcomes  $Y^t$ ,  $t \in \mathcal{T}$ , are expressible as

$$Y^t = \frac{1}{2}X^1 + X^2 + (2 \cdot \mathbb{1}[t=1] - 1)\frac{1}{4}X^1 + (2 \cdot \mathbb{1}[t=2] - 1)\frac{1}{4}X^2 + \epsilon_t,$$

where  $\epsilon_t \sim \mathcal{N}(0, \sigma)$ ,  $t \in \mathcal{T}$ , are i.i.d. noise terms. In our experiments, we will vary  $\sigma$  in the set  $\{0.1, 0.5, 0.8, 1.15, 1.5\}$  to study the impact on our method of increasing noise. Indeed, increasing the standard deviation of the noise impacting  $Y^t$ ,  $t \in \mathcal{T}$ , is expected to affect the quality of the estimates of  $m^t$  and  $\mu^t$  used by our sample-based dual-price queuing policy. From the above, in our notation of conditional mean treatment outcomes, we see that


$$m^0(X) = \frac{1}{4}X^1 + \frac{3}{4}X^2, \quad m^1(X) = \frac{3}{4}X^1 + \frac{3}{4}X^2, \quad \text{and} \quad m^2(X) = \frac{1}{4}X^1 + \frac{5}{4}X^2.$$

The historical/training set treatment assignments  $T$  come from a discrete distribution that, in line with our assumptions, depends on the covariates  $X$  only, i.e.,  $p(X) = [p^0(X), p^1(X), p^2(X)]$ , where  $p^t(x)$  is the assignment probability of treatment  $t$  given  $X = x$ . In particular, we let treatment assignment probabilities depend only on the treatment that is the best in expectation. This partitions the covariate space into three disjoint regions, each with associated values for  $p(x)$ : (1) when treatment 0 is the best (or at least as good) in expectation amongst the treatments (i.e.,  $0 = \min \arg \max_{t \in \mathcal{T}} m^t(x)$ ), we set  $p(x) = [0.8, 0.1, 0.1]$ ; (2) when treatment 1 is the best (or at least as good) in expectation amongst the treatments (i.e.,  $1 = \min \arg \max_{t \in \mathcal{T}} m^t(x)$ ), we set  $p(x) = [0.6, 0.3, 0.1]$ ; and (3) when treatment 2 is the best (or at least as good) in expectation amongst the treatments (i.e.,  $2 = \min \arg \max_{t \in \mathcal{T}} m^t(x)$ ), we set  $p(x) = [0.6, 0.1, 0.3]$ . This modeling choice mimics realistic scenarios where we anticipate that policy-/decision-makers frequently assign treatments “correctly.” In addition, for cases (2) and (3) where treatments 1 and 2 are the best in expectation and there is a higher probability of receiving those treatments relative to the probabilities in case (1), treatment 0 still has the highest probability at 60%. This modeling choice mimics realistic scenarios where there is a default no-treatment scenario (treatment 0) and there are limited quantities of treatments 1 and 2. Finally, we assume that the treatment capacities are known and given by  $b^0 = 1$ ,  $b^1 = 0.1$ , and  $b^2 = 0.05$ .

To evaluate the performance of our method on this data, we run 25 simulations. Each of these involves 360,000 test set samples (individuals) as a proxy for the asymptotic setting, and a number of training samples that we vary from 500 to 220,000.

**5.1.2. Choice of Outcome Estimators.** Since the asymptotic optimality of our policy  $\pi$  depends on the uniform almost sure convergence of  $\hat{m}^t$  to  $m^t$  (see Assumption 3 (ii)), we study the robustness of our method to Assumption 3 (ii) by using a variety of estimators for learning the mean treatment outcome models. If an incorrect parametric model is employed when learning  $m^t$ ,  $t \in \mathcal{T}$ , then such estimation errors may cause suboptimal policy learning. From the data generation process, we know that the true  $m^t$  functions are linear, see Section 5.1.1. Thus, we can use decision trees and lasso regression (with regularization parameter  $\alpha = 0.6$  so that lasso is misspecified) as our misspecified models (Assumption 3 (ii) does not hold) and  $k$ -NN as a non-parametric model with asymptotic consistency. Finally, we investigate if causal inference methods for ameliorating functional form bias in treatment effects estimation can improve policy learning in the presence of estimation errors of treatment outcomes. In particular, we use inverse propensity score weighted (IPW) and doubly robust (DR) variations of our misspecified models (decision trees and lasso regression) to see if they can mitigate bias from estimation errors. Since IPW and DR require estimating propensity scores, we use decision trees and logistic regression as the model classes for propensity score estimation. For an introduction to these methods, please see Battocchi et al. (2019).

**5.1.3. Performance and Discussion.** Our main results are summarized in Figure 1, where we plot our main performance metric under policies using various estimators of treatment outcomes. The figure shows our policy performance metric for each estimator in dependence of training sample size, with the left (resp. right) figure focused on the low (resp. high) variance case where  $\sigma = 0.1$  (resp.  $\sigma = 1.5$ ). For brevity, in the rest of this section we refer to various dual-price queuing policies using different estimators of the conditional mean treatment outcomes by just the estimator name. For example, the linear, lasso, decision tree policies in Figure 1 are the dual-price queuing policies based on ‘directly’ estimating the conditional mean outcomes using linear, lasso, and decision tree models, where linear is the correct model form while decision trees and lasso are misspecified models. The policies that start with ‘IPW’ and ‘DoublyRobust’ in their names are based on estimators of treatment outcomes that attempt to correct the bias of the lasso based policy by IPW or DR estimates of treatment outcomes. Finally, we have omitted from these results the performance of a ‘random’ policy that randomly assigns treatments as they arrive to waitlisted individuals (which performed near 0% irrespective of training sample size). Additional results for  $\sigma$  values of  $\{0.5, 0.8, 1.15\}$  can be found in Electronic Companion EC.4.1.



**Figure 1** Companion figure to the synthetic data results for  $\sigma = 0.1$  and  $\sigma = 1.5$ . Both subfigures show the ratio of out-of-sample performance of the sample-based queuing policy to that of the perfect foresight policy in dependence of training set size. Each line corresponds to a policy using a different treatment outcome estimator.

From Figure 1 (left), we see that, in the low variance setting, policy performance converges relatively fast (with under 9,000 training samples). This is expected: since observed outcomes affect

the estimation of counterfactuals, which in turn affect the estimation of policy parameters  $\hat{\mu}^*$ , low noise in the observed outcomes leads to faster convergence of policy performance. As we would expect, the linear estimator policy attains close to perfect performance since it correctly estimates the true conditional mean treatment outcomes. On the other hand, the various modified lasso policies appear to mitigate the estimation bias of the direct lasso estimates and improve policy performance. However, these results also show that the choice of propensity score estimates used to modify the direct treatment outcome estimates can impact policy performance. We see that policies using decision tree estimated propensity scores resulted in better convergence of policy performance (close to the linear policy performance) compared to policies using logistic regression estimated propensity scores. The worse performance of policies using logistic regression estimated propensity scores, compared to that of policies using decision tree estimated propensity scores, may result from the non-linear relation between treatment assignment probabilities and covariates. We emphasize that selecting an accurate propensity score model (even when the historical treatment assignment policy is unknown) is usually feasible in practice with modern machine learning approaches. In this low noise setting, non-parametric estimator based policies such as  $k$ -NN can effectively estimate the counterfactuals well enough to achieve close to perfect performance. Thus, as expected, in the direct estimator case, our method’s performance converges to 1 if the model employed is correct. On the other hand, in policies using IPW or DR adjustments, the choice of propensity score model and estimates impacts the convergence of policy performance. With a nonparametric estimator, convergence to 1 also occurs.

From Figure 1 (right), we see that in the high variance setting, the policies based on the non-parametric estimators of  $m^t$  all have a deterioration in performance relative to the low variance setting. In particular, the policy using  $k$ -NN estimators exhibits slower performance convergence as the high variance noise terms causes the  $k$ -NN mean outcome estimation error to decrease at a slower rate. On the other hand, compared to the low variance setting, the simpler parametric lasso model maintains a similar, though suboptimal, performance. The lasso model is more ‘robust’, in some sense, to more noise as both its mean treatment outcome estimation error and overall policy performance do not deteriorate much from the low to high variance setting. Finally, we see similar behavior in terms of improvement in policy performance from the IPW and DR correction methods to the lasso policy, though the degree of bias correction of IPW and DR methods depend on the estimated propensity scores used. These simple simulation results highlight the importance of model specification for counterfactual estimation and impact on model performance. We see that when Assumption 3 (ii) does not hold, allocation policies based on misspecified estimators generally do not converge to 100%. We also see that reweighting or doubly robust estimation of treatment outcomes can mitigate potential estimation bias. Finally, in Electronic Companion EC.4.2, we



investigate if our takeaways so far still hold in the nonlinear setting. We conduct an additional synthetic experiment with a similar setup but where replace the conditional mean treatment outcome functions with nonlinear quadratic functions. We find that the same qualitative takeaways still hold to a large degree.

## 5.2. Real Data Experiments: Allocating PSH and RRH Housing in LA

We next showcase the performance of our methodology on real data to design policies for allocating scarce resources to individuals experiencing homelessness in LA. We obtained our data from LAHSA through a Data Use Agreement and our data comes from the LA County Homeless Management Information System (HMIS) database (LAHSA 2015) and administered VI-SPDAT assessments (LAHSA 2017). We obtained Institutional Review Board approval for our data analysis. Additional details on the data sources and preprocessing can be found in Electronic Companion EC.4.3.

**5.2.1. HMIS and VI-SPDAT Data Description.** The HMIS data contains deanonymized data of individuals experiencing homelessness and their interactions with the system, which we refer to as *enrollments*. Each enrollment represents an instance in which an individual is assigned a specific resource, where the resource types (introduced in Section 1) are PSH, RRH, and S0. Note that S0 is a broad category that includes various services that are not considered permanent housing such as street outreach, homelessness prevention, emergency shelter, supportive services like childcare, and more. We also further breakout PSH into two sub-types PSH **Tenant-Based** and PSH **Site-Based** that differ, broadly speaking, in terms of supportive services available on-site and process of actually receiving housing. Further details can be found in in Electronic Companion EC.4.3. Since an individual can have multiple enrollments, we observe for each individual and enrollment pair the relevant date, and assigned resource type. The administered VI-SPDAT surveys, which are used for assessing individual vulnerability and resource prioritization, contain individual covariates, such as disabilities and prior housing history at the time of assessment.

We include all individuals with an assessment between the time period of 1/12/2015 (roughly when VI-SPDAT started being administered) to 12/31/2019 (to avoid idiosyncratic effects of Covid-19 during 2020) for a total of 63,764 samples. We use all individuals assessed between 1/12/2015 to 12/31/2017 as our training set to learn treatment outcomes and construct our policy, and then evaluate on individuals assessed between 1/1/2018 to 12/31/2019 to measure out-of-sample performance. We have a total of 4 possible resources, i.e.,  $\mathcal{T} = \{0, 1, 2, 3\}$ , which represent **no-treatment**, RRH, PSH **Tenant-Based**, and PSH **Site-Based**. To solve problem  $(\hat{\mathcal{D}})$  and derive a sample-based queuing policy, we use the training set capacities for **no-treatment**, RRH, PSH **Tenant-Based**, and PSH **Site-Based** as the estimates of capacity per person, which are 100%, 14.2%, 4.1%, and 4.6%, respectively. Since racial equality and equity are important concerns in our motivating example

in Section 1, we use race as the protected characteristic for evaluating allocation and outcome fairness. The protected features set is  $\mathcal{G} = \{\text{BlackAfAmerican}, \text{Hispanic}, \text{Other}, \text{White}\}$ , where **BlackAfAmerican** stands for Black or African American, and **Other** includes all other individuals who identify as American Indian or Alaska Native, Asian, Native Hawaiian or Other Pacific Islander, more than one racial group, or answered ‘Doesn’t know’. In this example, we consider non-White individuals to be minority groups that policymakers may want to prioritize due to historical discrimination.

*Outcome Definition.* According to the U.S. Department of Housing and Urban Development (HUD), the LA CoC should develop “action steps to end homelessness and prevent a return to homelessness” (LAHSA 2023). In line with this goal, we will use an individual’s return to homelessness after an intervention as the outcome to measure the performance of a policy. This outcome will serve as an input to our optimization model for designing a dual-price policy for allocating scarce housing resources effectively. Specifically, we focus on returns to homelessness within a two-year window following intervention for our real data experiments. However ‘returns to homelessness’ for individuals are not explicitly tracked. Instead, we construct a proxy outcome variable where an individual ‘returned to homelessness’ following intervention if we observe a subsequent enrollment into ‘emergency shelter’, ‘safe haven’, or ‘street outreach’. We chose these enrollment categories based on discussions with domain experts like matchers and case managers working within the LA CES since these types of interactions with the system indicate a need for homeless services and suggest an individual returned to homelessness again. For individuals without a treatment enrollment into **RRH** or **PSH**, we treat their first enrollment into a **S0** resource as the start of a **no-treatment** ‘intervention’. Therefore, we define a positive outcome,  $Y = 1$ , as *not* observing a subsequent return to homelessness within a two-year window of receiving intervention, and  $Y = 0$  otherwise. While other observation window lengths could be chosen, there is a trade-off in terms of window length. Shorter windows are potentially biased proxies since there is not enough time to observe post-intervention outcomes while longer windows result in less individuals for whom we can observe the entire window in our dataset.

*Policy Evaluation.* We learn three different versions of our proposed policy from the training set and evaluate their performance by their overall proportion of positive outcomes as the ‘effectiveness’ metric and their fairness properties in terms of statistical parity in allocation (4) and in outcomes (7) on the test set. We emphasize that our evaluation metric is a counterfactual probability, since the true counterfactuals under assignments different from the historical ones are not available. We explain how we generate these counterfactuals, as well as how we evaluate out-of-sample performance, in Section 5.2.2. To investigate the fairness properties, we look at discrepancies in the percentage receiving housing resources and discrepancies in the proportion

of positive outcomes between racial groups in the out-of-sample allocation. In our results, **Base** refers to the policy without fairness constraints, while the other two refer to fairness-constrained variations defined in Section 4. Based on our motivating example, we chose to focus on fairness extensions that prioritized minority groups in terms of either allocation, **Alloc Min Priority**, or outcomes, **Outcome Min Priority**. To benchmark performance, we compare the three policies to the **Historical** policy, which uses the same assignments originally found in the data, in terms of the overall proportion of positive outcomes and the two statistical parity metrics. In this way, we can see if our proposed policies can improve on the ‘effectiveness’ and fairness of the currently deployed policy. Additionally, we include the performance of a **VI-SPDAT Thresholding** policy, which learns, from historical data, VI-SPDAT score thresholds for allocation so that higher vulnerability scores get more supportive resources like PSH. These thresholds are then used online to assign resources based on an individual’s VI-SPDAT score. This offers a comparison point of how the VI-SPDAT score can be used for prioritization. Finally, we also compare the performance of overall outcomes to a **Perfect Foresight** policy that implements the sample-based dual-price queuing policy while knowing both the test set distribution of arriving individuals and their counterfactuals.

**5.2.2. Counterfactuals and Outcome Estimators.** Unlike the synthetic experiment, we do not have access to counterfactuals for each individual under treatments other than the one received. Therefore, we take a semi-synthetic approach and use model generated counterfactuals for the test set by choosing models that are well calibrated to the entire data, therefore including data not used by the training estimators. We choose calibration as the measure of fit since we are evaluating based on expected outcomes of a binary variable, and therefore we want our predicted probabilities of a positive outcome to match the observed data. In particular, to correct for selection bias in our observational data, we use weighted calibration for model validation, where we calculate average positive outcomes using the (self-normalized/Hájek) IPW estimator instead of simple averaging of observed outcomes (Hernan and Robins 2023).

To fit counterfactual models, we consider the three causal estimation approaches of direct, IPW, and DR that were also used in our synthetic experiment as discussed in Section 5.1.2. As some of these methods may involve estimating both the observed outcomes and propensity scores, our choice of base estimators for the outcomes and propensities include logistic regression (with and without regularization), random forests, gradient boosted trees (Friedman 2001), and XGBoost trees (Chen and Guestrin 2016). For model selection, we use standard cross-validation to find the best hyperparameters for each model class that minimize log-loss for predicting outcomes or the assigned treatment for propensity scores and then select the final model based on weighted calibration curves on a held-out validation set.

An additional concern of our model-generated counterfactuals is potential bias by race because our evaluation is based on these model predicted probabilities. If the outcome models are miscalibrated for a particular racial group (e.g., they systematically over- or under-estimates a group’s true positive outcome probabilities under a treatment), then the overall policy outcome numbers we report will have those errors and unfairly favor or penalize a group. Though this was not part of our model selection process, we investigated potential racial bias by comparing calibration curves for each racial group across treatments using our selected counterfactual-generating models. For the most part, the calibration curves across racial groups are similar except for the ‘Other’ racial group, likely because this group contained too few samples to get accurate estimates. For additional details, please see Section EC.4.3.3.

While the above addresses generating counterfactuals for evaluation, we take a similar approach for actually learning  $\hat{m}^t$  functions but *only* use the training data. For more details on counterfactual generation, model calibration, and racial fairness of the estimates, please see the Electronic Companions EC.4.3.2 and EC.4.3.3.

**5.2.3. Experimental Results.** Our efficiency results are summarized in Table 1, which shows the performance of the proposed policies by their out-of-sample expected outcomes and their relative percentage difference compared to the historical policy (e.g., the relative percentage change of **Base** versus **Historical** is calculated as the outcome of **Base** minus the outcome of **Historical**, divided by the outcome of **Historical**). Since the **VI-SPDAT Thresholding** performance is largely similar to **Historical**, any comparison to the **Historical** policy also applies to the **VI-SPDAT Thresholding** policy. As a point of reference for the effectiveness of different allocations, we also included the potential population expected outcome if every individual receives **no-treatment**.

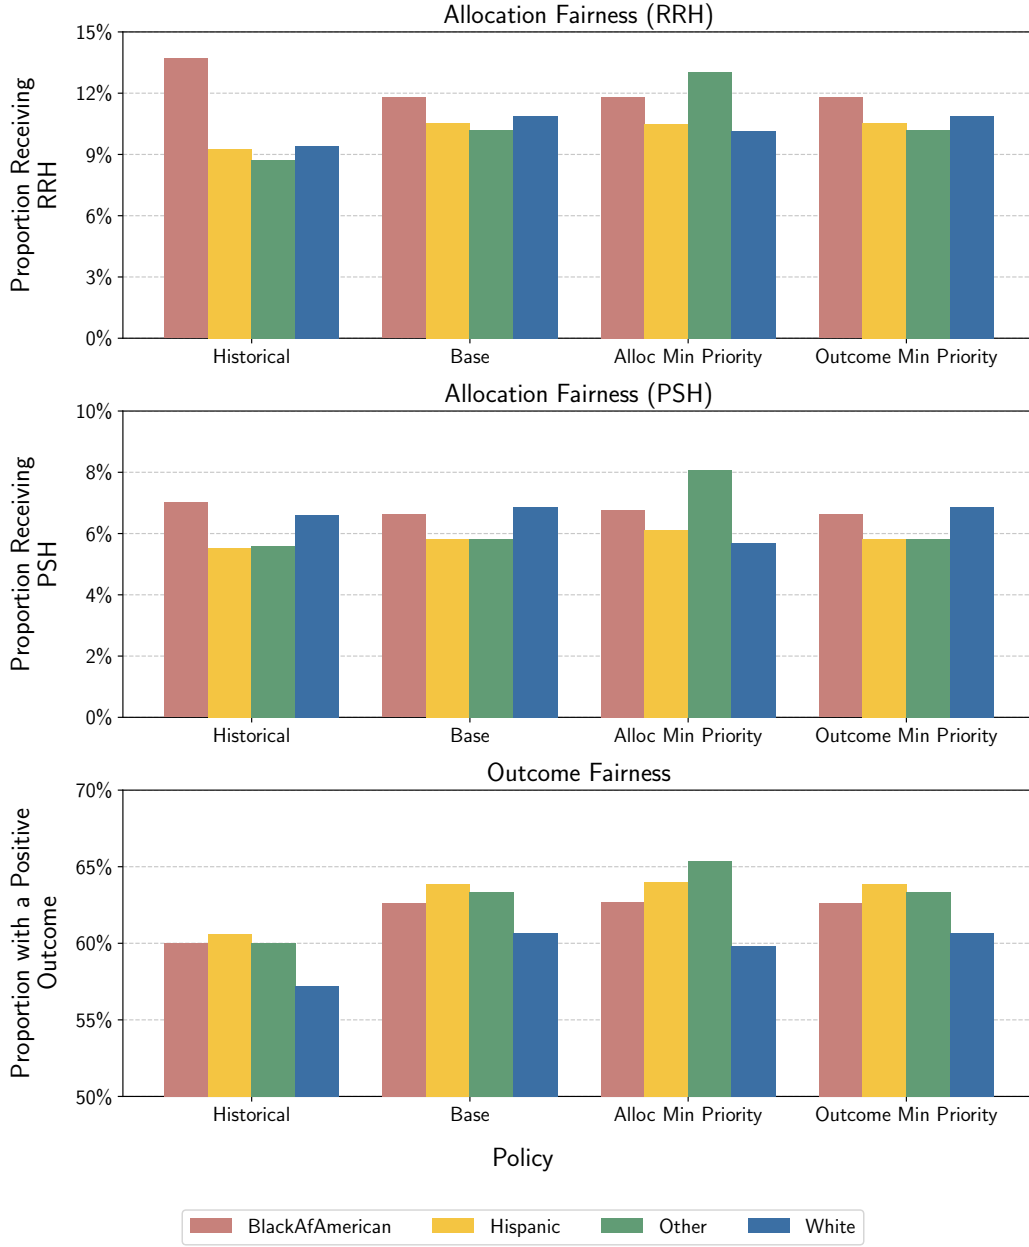
In terms of the out-of-sample expected outcomes, we see in Table 1 that our base model achieves a 5.16% improvement above the **Historical** policy, or an absolute 3.07 percentage point (p.p.) increase, while having a performance gap of 1.55%, or 0.92 p.p., compared to the **Perfect Foresight** policy. We found that the performance gap relative to **Perfect Foresight** comes from two sources of error: (i) estimation error of treatment outcomes  $\hat{m}^t$ , which are plug-ins into the assignment policy, and (ii) shifts in the distribution of individuals we observe and capacities per person for each treatment from training to testing, which affects the optimal value of  $\mu^*$  and any estimates  $\hat{\mu}^*$  of it. These sources of error could come from insufficient training samples for estimating outcomes or observing the distribution of individuals, or from violations of assumptions needed such as stationary distribution from training to testing, or positivity assumption for causal inference. Practically speaking, we see that the **Historical** policy already performs rather well compared to the **Perfect Foresight** policy given the scarce amount of resources. Our policy further improves upon **Historical**, leading to roughly 500 more individuals exiting homelessness per

year. This shows our proposed policy can be a *systematic* method to assist and supplement the experience and knowledge of domain experts.

**Table 1** Companion table to the real data experiment showing out-of-sample proportion of positive outcomes under each policy, and the relative percentage difference compared to the Historical allocation.

Policy	Proportion of Positive Outcomes	Percentage Change vs Historical
No Treatment	54.25%	−8.76%
Historical	59.46%	0.00%
VI-SPDAT Thresholding	59.80%	0.01%
Base	62.53%	5.16%
Alloc Min Priority	62.51%	5.13%
Outcome Min Priority	62.53%	5.16%
Perfect Foresight	63.45%	6.71%

Figure 2 illustrates the allocation fairness of each policy by showing the proportion of each racial group receiving RRH and PSH and the outcome fairness of each policy by showing the expected outcome within each racial group. While **Base** improves upon the expected outcomes of **Historical**, **Base** has marginally worse allocation fairness properties than **Historical** when considering the proportion of minority groups receiving PSH and has more uniform allocation rates across all groups for RRH, as displayed in Figure 2. However, by adding appropriate fairness constraints, the **Alloc Min Priority** policy is able to mitigate any problems and achieve better prioritization of minority groups for allocations. The first subplot of Figure 2 shows the proportion of each racial group receiving RRH under each policy. We see that the **Base** and **Outcome Min Priority** policies lead to more uniform allocation rates across groups but still have RRH allocation disparity gaps between **White** vs **Hispanic** and **Other** groups compared to the **Historical**. **Alloc Min Priority**, however, satisfies the minority allocation prioritization fairness constraint presented in Section 4.1 out-of-sample. Looking at PSH allocations in the second subplot of Figure 2, we see that the **Historical** policy does not prioritize the **Hispanic** and **Other** racial groups to be at least as likely to receive PSH as the majority group. Our **Base** and **Outcome Min Priority** policies maintain similar allocation disparity gaps as **Historical** while also allocating slightly less now to **BlackAfAmerican** individuals relative to **White** individuals. However, if we explicitly add allocation fairness constraints to our sample-based problem to prioritize minority groups for receiving PSH as shown in Section 4.1, then our **Alloc Min Priority** is able to achieve our desired fairness goal.



**Figure 2** Companion figure to the real data experiments: Out-of-sample fairness results showing the proportion of each racial group receiving RRH, PSH (Allocation Fairness) and the proportion of each racial group with a positive outcome (Outcome Fairness).

In the last subplot of Figure 2, we show the proportion of positive outcomes within each racial group for each policy and focus on the prioritization of minority groups to have groupwise outcomes that are at least as good as the majority group. In this case, the **Historical** policy already achieves this notion of fairness and **Base** does not create outcome disparities between minority groups compared to the majority group. The **Outcome Min Priority** policy, which explicitly includes outcome fairness constraints in the sample problem, also achieves the desired fairness goal while also improv-

ing **BlackAfAmerican**, **Hispanic**, and **White** individual outcomes relative to **Historical** (this is also true for **Base**). We also see in these results that **Alloc Min Priority** has favorable fairness properties by achieving both fairness notions of prioritization of minority groups for allocation and outcomes, though this need not hold in other problems. In general, allocation statistical parity based fairness notions and constraints are incompatible with outcome statistical parity based fairness (Jo et al. 2023). Therefore, depending on the desired fairness goals, policymakers may need to choose between alternative fairness notions, or choose a combination of them to achieve a compromise.

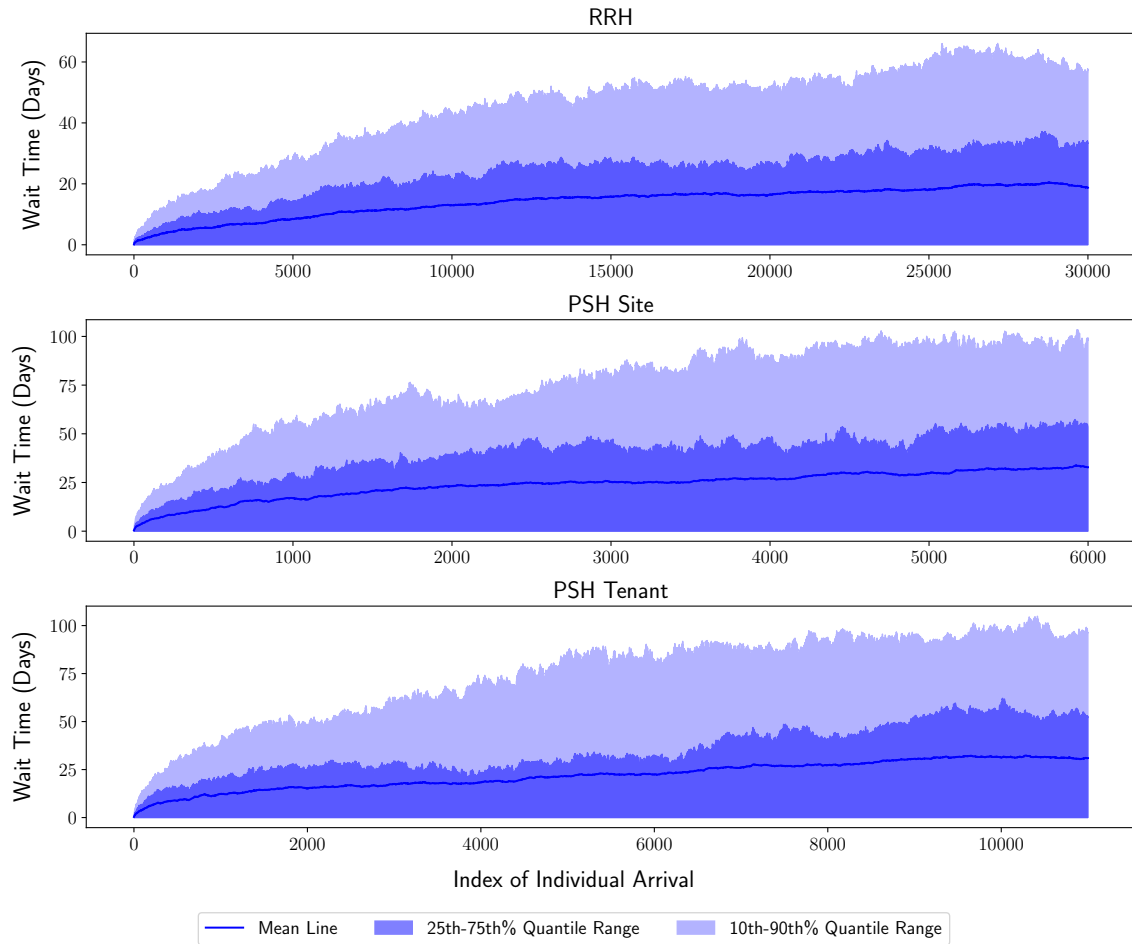
From Figure 2, we can conclude that fairness-constrained policies such as **Alloc Min Priority** and **Outcome Min Priority** can improve the fairness properties of the **Base** policy and satisfy the fairness constraints specified in Section 4. We note that although our constraints are imposed at the statistical level, they did hold for the sample path that was realized in our test data. More interestingly, we see in Table 1 that our fairness-constrained policies suffer almost no performance gap in terms of expected outcomes compared to **Base**, suggesting almost no ‘price of fairness’. Intuitively, we would expect some loss of performance by enforcing certain constraints on allocations. This may either be due to this application or the flexibility of our methodology to find well-performing policies that satisfy various fairness goals.

Beyond our quantitative takeaways above, we additionally analyze the qualitative behavior of our policy in the Electronic Companion EC.4.4. One potential concern for our policy is that by optimizing for overall expected outcomes, it may overlook vulnerable individuals who are not responsive to treatments. However, we find in Section EC.4.4.1 that this concern does not occur in our experimental results because the most vulnerable individuals also tend to benefit the most from treatment and therefore are prioritized by our policy. Additionally, policymakers may want to understand why our policy allocates resources differently from the existing **Historical** policy. In Section EC.4.4.2, we analyze cases where our policy differs from the **Historical** based on individual covariates. We find that allocation differences were due to covariates that did not come directly from the VI-SPDAT survey but from auxiliary HMIS data sources, such as an individual’s current employment type/status. Please see the Electronic Companion EC.4.4 for additional details.

**5.2.4. Queue Wait Times.** Since our dual-price queuing policy involves individuals waiting until they are at the front of the queue and the associated resource becomes available, it is important to understand the wait time properties of our policy and how it affects individuals. We clarify that the waiting time for matching *prior* to taking the VI-SPDAT is an explicit input to our model and here we evaluate the wait time after the assessment. To that end, we additionally evaluate the distribution of individual waiting times within each treatment queue under our **Base** policy

and the impact of increasing wait times on individual outcomes. To do so, we run 250 simulations of individual and resource arrivals over a 10,000 day horizon to evaluate the long-run result (our original test period in Section 5.2.3 was 730 days). We simulate resource arrivals using a Poisson arrival process, where the exponential inter-arrival times are parameterized by the average historical resource arrival rates per day. We similarly simulate individual arrivals using a Poisson arrival process. For each arriving individual, we sample their covariates uniformly with replacement from the historical population.

In Figure 3, we show wait time statistics and their progression as more individuals arrive in the system. Within each sample path, for each individual assigned to RRH, PSH Site, or PSH Tenant queues, we compute their wait time as the time difference between their arrival time and the time they are matched with a resource. We compute the mean, 10<sup>th</sup>, 25<sup>th</sup>, 75<sup>th</sup>, and 90<sup>th</sup> percentile wait times of the  $i^{\text{th}}$  individual to arrive in the system across all 250 simulations.



**Figure 3** Distribution of individual wait times for different treatments across 250 simulated sample paths.



We see that the mean waiting time stabilizes at approximately 18, 30, and 30 days for RRH, PSH Site-Based, and PSH Tenant-Based, respectively, while 90<sup>th</sup> percentile wait times are 55, 100, and 100 days for RRH, PSH Site-Based, and PSH Tenant-Based, respectively. These wait times are much shorter than the historical wait times observed in our data, where the mean wait times from 2015 to 2019 were 213, 373, and 266 days for individuals assigned to RRH, PSH Site-Based, and PSH Tenant-Based, respectively. Finally, to understand how reasonable these wait times are, we investigate their effect on outcomes and find that, while the wait times negatively impact potential outcomes, the wait time levels in our numerical experiments result in relatively small decreases in expected conditional outcomes; see Electronic Companion EC.5.3 for additional details. For further discussions around wait time and details on the experiment setup and results, please see Electronic Companion EC.5.

## Disclaimer

The views and opinions of the data presented belong to the authors and do not represent the views or opinions of LAHSA.

## Acknowledgments

P. Vayanos is funded in part by the National Science Foundation under CAREER award number 2046230. B. Tang is funded in part by the National Science Foundation Graduate Research Fellowship Program (GRFP). E. Rice, P. Vayanos, and B. Tang are funded in part 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. The authors are grateful for the support.

## References

- Adelman D (2007) Dynamic bid prices in revenue management. *Operations Research* 55(4):647–661.
- Aggarwal G, Goel G, Karande C, Mehta A (2011) Online vertex-weighted bipartite matching and single-bid budgeted allocations. *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, 1253–1264.
- Ahani N, Andersson T, Martinello A, Teytelboym A, Trapp AC (2021) Placement optimization in refugee resettlement. *Operations Research* 69(5):1468–1486.
- Ahani N, Gözl P, Procaccia AD, Teytelboym A, Trapp AC (2024) Dynamic placement in refugee resettlement. *Operations Research* 72(3):1087–1104.
- Athanassoglou S, Sethuraman J (2011) House allocation with fractional endowments. *International Journal of Game Theory* 40:481–513.
- Athey S, Wager S (2021) Policy learning with observational data. *Econometrica* 89(1):133–161.

- Azizi MJ, Vayanos P, Wilder B, Rice E, Tambe M (2018) Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. *Integration of Constraint Programming, Artificial Intelligence, and Operations Research: 15th International Conference*, 35–51.
- Balseiro SR, Besbes O, Pizarro D (2024) Survey of dynamic resource-constrained reward collection problems: Unified model and analysis. *Operations Research* 72(5):2168–2189.
- Bansak K, Ferwerda J, Hainmueller J, Dillon A, Hangartner D, Lawrence D, Weinstein J (2018) Improving refugee integration through data-driven algorithmic assignment. *Science* 359(6373):325–329.
- Bansak K, Paulson E (2024) Outcome-driven dynamic refugee assignment with allocation balancing. *Operations Research* 72(6):2375–2390.
- Battocchi K, Dillon E, Hei M, Lewis G, Oka P, Oprescu M, Syrgkanis V (2019) EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation. <https://github.com/py-why/EconML>, version 0.15.1.
- Bertsekas D (2009) *Convex Optimization Theory*. Athena Scientific Optimization and Computation Series (Athena Scientific).
- Bertsimas D, Dunn J, Mundru N (2019) Optimal prescriptive trees. *INFORMS Journal on Optimization* 1(2):164–183.
- Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61(1):73–87.
- Bhattacharya D, Dupas P (2012) Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics* 167(1):168–196.
- Chen T, Guestrin C (2016) Xgboost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.
- Chen X, Christensen TM (2015) Optimal uniform convergence rates and asymptotic normality for series estimators under weak dependence and weak conditions. *Journal of Econometrics* 188(2):447–465.
- Cheng PE (1984) Strong consistency of nearest neighbor regression function estimators. *Journal of Multivariate Analysis* 15(1):63–72.
- Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W, Robins J (2018) Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21(1):C1–C68.
- Corbett-Davies S, Pierson E, Feller A, Goel S, Huq A (2017) Algorithmic decision making and the cost of fairness. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 797–806.
- Devanur NR, Jain K, Kleinberg RD (2013) Randomized primal-dual analysis of ranking for online bipartite matching. *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete algorithms*, 101–107.

- Dickerson J, Sandholm T (2015) Futurematch: Combining human value judgments and machine learning to match in dynamic environments. *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 622–628.
- Dudík M, Langford J, Li L (2011) Doubly robust policy evaluation and learning. *Proceedings of the 28th International Conference on Machine Learning*, 1097–1104.
- Dwork C, Hardt M, Pitassi T, Reingold O, Zemel R (2012) Fairness through awareness. *Proceedings of the 3rd innovations in theoretical computer science conference*, 214–226.
- Foster DJ, Syrgkanis V (2023) Orthogonal statistical learning. *The Annals of Statistics* 51(3):879–908.
- Freeman R, Shah N, Vaish R (2020) Best of both worlds: Ex-ante and ex-post fairness in resource allocation. *Proceedings of the 21st ACM Conference on Economics and Computation*, 21–22.
- Freund D, Lykouris T, Paulson E, Sturt B, Weng W (2023) Group fairness in dynamic refugee assignment. *Proceedings of the 24th ACM Conference on Economics and Computation*, 701–701.
- Friedman JH (2001) Greedy function approximation: a gradient boosting machine. *Annals of Statistics* 29(5):1189–1232.
- Gneiting T, Raftery AE (2007) Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477):359–378.
- Grand-Clément J, Chan CW, Goyal V, Escobar G (2023) Robustness of proactive intensive care unit transfer policies. *Operations Research* 71(5):1653–1688.
- Hansen BE (2004) Uniform convergence rates for nonparametric estimation.
- Hernan M, Robins J (2023) *Causal Inference: What If*. Boca Raton: Chapman & Hall (CRC Press).
- Jacob D (2021) Cate meets ml: Conditional average treatment effect and machine learning. *Digital Finance* 3(2):99–148.
- Jo N, Aghaei S, Gómez A, Vayanos P (2021) Learning optimal prescriptive trees from observational data. *arXiv preprint arXiv:2108.13628* .
- Jo N, Tang B, Dullerud K, Aghaei S, Rice E, Vayanos P (2023) Fairness in contextual resource allocation systems: Metrics and incompatibility results. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11837–11846.
- Johnston CM, Bleszenohl S, Vayanos P (2020) Preference elicitation and aggregation to aid with patient triage during the covid-19 pandemic. *Workshop on Participatory Approaches to Machine Learning*.
- Kallus N (2017) Recursive partitioning for personalization using observational data. *Proceedings of the 34th International Conference on Machine Learning*, 1789–1798.
- Karp RM, Vazirani UV, Vazirani VV (1990) An optimal algorithm for on-line bipartite matching. *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, 352–358.

- Kaya YB, Maass KL, Dimas GL, Konrad R, Trapp AC, Dank M (2024) Improving access to housing and supportive services for runaway and homeless youth: Reducing vulnerability to human trafficking in new york city. *IISE Transactions* 56(3):296–310.
- Kennedy EH (2023) Towards optimal doubly robust estimation of heterogeneous causal effects. *Electronic Journal of Statistics* 17(2):3008–3049.
- Kitagawa T, Tetenov A (2018) Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica* 86(2):591–616.
- Kube A, Das S, Fowler PJ (2019) Allocating interventions based on predicted outcomes: A case study on homelessness services. *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, volume 33, 622–629.
- Kube AR, Das S, Fowler PJ (2023) Fair and efficient allocation of scarce resources based on predicted outcomes: implications for homeless service delivery. *Journal of Artificial Intelligence Research* 76:1219–1245.
- Kunnumkal S, Talluri K (2016) On a piecewise-linear approximation for network revenue management. *Mathematics of Operations Research* 41(1):72–91.
- Künzel SR, Sekhon JS, Bickel PJ, Yu B (2019) Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences* 116(10):4156–4165.
- Li D, Pang Z, Qian L (2023) Bid price controls for car rental network revenue management. *Production and Operations Management* 32(1):261–282.
- Li Z, Lieberman K, Macke W, Carrillo S, Ho CJ, Wellen J, Das S (2019) Incorporating compatible pairs in kidney exchange: A dynamic weighted matching model. *Proceedings of the 20th ACM Conference on Economics and Computation*, 349–367.
- Liero H (1989) Strong uniform consistency of nonparametric regression function estimates. *Probability Theory and Related Fields* 82(4):587–614.
- Los Angeles Homeless Services Authority (2015) LA HMIS Policies and Procedures. <https://www.lahsa.org/documents?id=1128-la-hmis-policies-and-procedures.pdf&ref=hmis>.
- Los Angeles Homeless Services Authority (2017) CES Survey for Individuals Survey Packet. <https://www.lahsa.org/documents?id=1306-form-1306-ces-survey-for-individuals-survey-packet.pdf>.
- Los Angeles Homeless Services Authority (2018a) An Introduction to the Coordinated Entry System & How to Conduct the CES Triage Tools. <http://publichealth.lacounty.gov/sapc/docs/providers/special-programs/Coordinated%20Entry%20HMIS%20Training.pdf>.
- Los Angeles Homeless Services Authority (2018b) Report and Recommendations of the Ad Hoc Committee on Black People Experiencing Homelessness. <https://www.lahsa.org/documents?id=2823-report-and-recommendations-of-the-ad-hoc-committee-on-black-people-experiencing-homelessness.pdf>.

- Los Angeles Homeless Services Authority (2023) Los Angeles Continuum of Care. <https://www.lahsa.org/coc/>.
- Los Angeles Homeless Services Authority (2024a) 2024 Greater Los Angeles Homeless Count. <https://www.lahsa.org/documents?id=8164-2024-greater-los-angeles-homeless-count-results-long-version-.pdf>.
- Los Angeles Homeless Services Authority (2024b) 2024 Housing Inventory Count. <https://www.lahsa.org/documents?id=8162-2024-housing-inventory-count.xlsx>.
- Mack Yp, Silverman BW (1982) Weak and strong uniform consistency of kernel regression estimates. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 61:405–415.
- Manshadi V, Niazadeh R, Rodilitz S (2021) Fair dynamic rationing. *Proceedings of the 22nd ACM Conference on Economics and Computation*, 694–695.
- Mashiat T, Gitiaux X, Rangwala H, Fowler P, Das S (2022) Trade-offs between group fairness metrics in societal resource allocation. *Proceedings of the 5th ACM Conference on Fairness, Accountability, and Transparency*, 1095–1105.
- Mehta A, Saberi A, Vazirani U, Vazirani V (2007) Adwords and generalized online matching. *Journal of the ACM* 54(5):22–es.
- Mehta A, et al. (2013) Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science* 8(4):265–368.
- Milburn NG, Edwards E, Obermark D, Rountree J (2021) Inequity in the permanent supportive housing system in los angeles: Scale, scope and reasons for black residents’ returns to homelessness. <https://www.capolicylab.org/wp-content/uploads/2021/10/Inequity-in-the-PSH-System-in-Los-Angeles.pdf>.
- Nguyen Q, Das S, Garnett R (2021) Scarce societal resource allocation and the price of (local) justice. *Proceedings of the AAAI Conference on Artificial Intelligence*, 5628–5636.
- Nie X, Wager S (2021) Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika* 108(2):299–319.
- OrgCode (2015) Vulnerability index - service prioritization decision assistance tool (vi-spdatt): Prescreen triage tool for single adults. <https://www.havenforhope.org/wp-content/uploads/2018/11/VI-SPDAT-v2.01-Family-US-Fillable.pdf>.
- OrgCode (2020) The Time Seems Right: Let’s Begin the End of the VI-SPDAT. <https://www.orgcode.com/blog/the-time-seems-right-lets-begin-the-end-of-the-vi-spdatt>.
- OrgCode Consulting, Inc and Community Solutions (2015) Changes in the VI-SPDAT v2.0. [https://www.ncceh.org/media/files/page/359d0ab6/VI-SPDAT\\_v2\\_Backgrounder.pdf](https://www.ncceh.org/media/files/page/359d0ab6/VI-SPDAT_v2_Backgrounder.pdf).
- Qian M, Murphy SA (2011) Performance guarantees for individualized treatment rules. *Annals of Statistics* 39(2):1180.

- Radovanovic D, Pini S, Franceschi E, Pecis M, Airoidi A, Rizzi M, Santus P (2021) Characteristics and outcomes in hospitalized covid-19 patients during the first 28 days of the spring and autumn pandemic waves in milan: an observational prospective study. *Respiratory Medicine* 178:106323.
- Rahmattalabi A, Jabbari S, Lakkaraju H, Vayanos P, Izenberg M, Brown R, Rice E, Tambe M (2021) Fair influence maximization: a welfare optimization approach. *Proceedings of the AAAI Conference on Artificial Intelligence*, 11630–11638.
- Rahmattalabi A, Vayanos P, Dullerud K, Rice E (2022) Learning resource allocation policies from observational data with an application to homeless services delivery. *Proceedings of the 5th ACM Conference on Fairness, Accountability, and Transparency*, 1240–1256.
- Shapiro A, Dentcheva D, Ruszczyński A (2014) *Lectures on Stochastic Programming: Modeling and Theory, Second Edition* (SIAM).
- Swaminathan A, Joachims T (2015) Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research* 16(1):1731–1755.
- Talluri K, Van Ryzin G (1998) An analysis of bid-price controls for network revenue management. *Management Science* 44(11-part-1):1577–1593.
- Talluri KT, Van Ryzin G (2004) *The Theory and Practice of Revenue Management*, volume 1 (Springer).
- US Department of Housing and Urban Development (2021) FY 2022 HMIS Data Standards Data Dictionary. <https://files.hudexchange.info/resources/documents/HMIS-Data-Dictionary.pdf>.
- US Department of Housing and Urban Development Office of Policy Development and Research (2007) The Applicability of Housing First Models to Homeless Persons with Serious Mental Illness. <https://www.huduser.gov/portal/publications/hsgfirst.pdf>.
- Vera A, Banerjee S (2021) The bayesian prophet: A low-regret framework for online decision making. *Management Science* 67(3):1368–1391, URL <http://dx.doi.org/10.1287/mnsc.2020.3624>.
- Xie H (2023) Uniform convergence results for the local linear regression estimation of the conditional distribution. *Statistics & Probability Letters* 201:109886.
- Zhang D, Weatherford L (2017) Dynamic pricing for network revenue management: A new approach and application in the hotel industry. *INFORMS Journal on Computing* 29(1):18–35.
- Zhou Z, Athey S, Wager S (2023) Offline multi-action policy learning: Generalization and optimization. *Operations Research* 71(1):148–183.

## Online Appendix

### EC.1. Dual-Price Queuing Policy under Full Distributional Information

In this section, we assume that we know the joint distribution of the covariates  $X$  and the potential outcomes  $Y^t$ ,  $t \in \mathcal{T}$ . We will characterize a dual-price queuing policy under this assumption and show that is optimal in problem  $(\mathcal{P})$ . This characterization and result will allow us to establish the asymptotic optimality of the sample-based dual-price queuing policy in Theorem 1.

To this end, we consider problem  $(\mathcal{P})$  and its Lagrangian dual

$$\nu^* = \min_{\mu \in \mathbb{R}_+^{m+1}} \max_{\pi \in \Pi} \mathbb{E} \left[ \sum_{t \in \mathcal{T}} \pi^t(X) (m^t(X) - \mu^t) \right] + \sum_{t \in \mathcal{T}} \mu^t b^t. \quad (\text{EC.1})$$

Similarly to the construction in Section 3, for any  $\mu \in \mathbb{R}_+^{m+1}$ , the inner maximization problem is solved by the policy

$$\pi^t(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} (m^{t'}(x) - \mu^{t'}) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T}, \forall x \in \mathcal{X}. \quad (\text{EC.2})$$

By substituting policy (EC.2) into the dual problem (EC.1), we obtain

$$\nu^* = \min_{\mu \in \mathbb{R}_+^{m+1}} \mathbb{E} \left[ \max_{t \in \mathcal{T}} (m^t(X) - \mu^t) \right] + \sum_{t \in \mathcal{T}} \mu^t b^t. \quad (\mathcal{D})$$

In the remainder, we denote by  $\mu^*$  an optimal solution to  $(\mathcal{D})$ , which exists because we can restrict the feasible set of  $(\mathcal{D})$  to a compact set without loss of generality; see Lemma EC.3. By choosing  $\mu$  as  $\mu^*$  in (EC.2), we define a candidate policy  $\pi^*$  through

$$\pi^{*,t}(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} (m^{t'}(x) - \mu^{*,t'}) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T}, \forall x \in \mathcal{X}. \quad (\text{EC.3})$$

Next, we introduce the queuing implementation of policy  $\pi^*$ .

**DEFINITION EC.1 (DUAL-PRICE QUEUING POLICY).** Establish  $m$  queues each associated with a treatment  $t \in \mathcal{T} \setminus \{0\}$ . At any time, individuals in queue  $t$  are waitlisted for treatment  $t$ . There are two events triggering an action: (i) arrival of an individual, and (ii) arrival of a treatment.

- (i) When an individual with covariate vector  $x$  arrives, compute the treatment  $t = \min \arg \max_{t' \in \mathcal{T}} m^{t'}(x) - \mu^{*,t'}$ , where  $t$  attains the maximum difference between the conditional mean treatment outcome and the associated dual variable. If  $t \neq 0$ , assign them to queue  $t$ , otherwise assign them no-treatment, meaning they will not wait for a particular treatment and will exit the system. If they are the only individual waiting in queue  $t$ , check whether there is an available treatment of type  $t$ . If there is, assign treatment  $t$  to the individual.

- (ii) When a treatment of type  $t$  becomes available, assign this treatment to the first individual in queue  $t$ . If there is no one in queue  $t$ , keep the treatment until an individual is assigned to this queue.

Similarly to the sample-based dual-price queuing policy, the dual-price queuing policy assigns  $t = 0$  to individuals still waiting in queues at the end of the time horizon when implemented in finite horizon. We note that in general,  $\mu^*$ , the optimal solution to  $(\mathcal{D})$ , may not be unique since the objective function of  $(\mathcal{D})$  may not be *strictly* convex; this depends on the form of the joint distribution of  $(X, \{Y^t\}_{t \in \mathcal{T}})$ . This means that the candidate policy  $\pi^*$  defined in (EC.3) and the dual-price queuing policy may not be uniquely specified.

The next theorem shows that the dual-price queuing policy satisfies the capacity constraints in problem  $(\mathcal{P})$ , and its asymptotic expected average outcome matches the optimal value  $z^*$  of  $(\mathcal{P})$ . It thus solves problem  $(\mathcal{P})$ .

**THEOREM EC.1.** *Under Assumption 2, the dual-price queuing policy is optimal in problem  $(\mathcal{P})$ .*

The proof of Theorem EC.1 relies on establishing some structural properties of problem  $(\mathcal{P})$  and applying standard techniques based on optimality conditions from network revenue management (Talluri and Van Ryzin 1998), and more generally, from capacity-constrained dynamic programming (Balseiro et al. 2024). The idea of the proof is that it is sufficient to show that  $\pi^*$  is feasible and optimal in  $(\mathcal{P})$ . Since strong duality holds for problem  $(\mathcal{D})$ , the optimal solutions of  $(\mathcal{D})$  and its dual must satisfy the Karush-Kuhn-Tucker (KKT) conditions. The KKT conditions imply that  $\pi^*$ , which is constructed from an optimal solution  $\mu^*$  of  $(\mathcal{D})$ , is a feasible solution of  $(\mathcal{P})$  and yields an objective value equal to  $v^* = z^*$ . Full proof can be found in the next appendix section.

In contrast to the sample-based dual-price queuing policy, the dual-price queuing policy in Definition EC.1 is unidentifiable because its definition relies on the joint distribution of  $(X, \{Y^t\}_{t \in \mathcal{T}})$ , which is unknown in reality.

## EC.2. Proofs

In this section, we provide the proofs of the results in Sections EC.1 and 3.

### EC.2.1. Proofs for Section EC.1

The proof of Theorem EC.1 requires the following two technical lemmas.

**LEMMA EC.1.** *For any  $\mu \in \mathbb{R}_+^{m+1}$ , denote by  $\mu^{-t}$  its subvector  $(\mu^0, \dots, \mu^{t-1}, \mu^{t+1}, \dots, \mu^m) \in \mathbb{R}_+^m$  excluding  $\mu^t$ , and define the random variable  $Z^{\mu^{-t}} = m^t(X) - \max_{t' \in \mathcal{T} \setminus \{t\}} (m^{t'}(X) - \mu^{t'})$  for all  $t \in \mathcal{T}$ . Under Assumption 2,  $Z^{\mu^{-t}}$  is continuously distributed for any  $t \in \mathcal{T}$  and  $\mu^{-t} \in \mathbb{R}_+^m$ .*



*Proof.* We will show that  $\mathbb{P}(Z^{\mu^{-t}} = z) = 0$  for all  $t \in \mathcal{T}$ ,  $\mu^{-t} \in \mathbb{R}_+^m$ , and  $z \in \mathbb{R}$ . To this end, fix arbitrary  $t \in \mathcal{T}$ ,  $\mu^{-t} \in \mathbb{R}_+^m$ , and  $z \in \mathbb{R}$ . We have  $\mathbb{P}(Z^{\mu^{-t}} = z) \geq 0$  so we only need to show that  $\mathbb{P}(Z^{\mu^{-t}} = z) \leq 0$ . We have

$$\begin{aligned} \mathbb{P}(Z^{\mu^{-t}} = z) &= \mathbb{P}(m^t(X) - \max_{t' \in \mathcal{T} \setminus \{t\}} (m^{t'}(X) - \mu^{t'}) = z) \\ &= \mathbb{P}(\min_{t' \in \mathcal{T} \setminus \{t\}} (m^t(X) - m^{t'}(X) + \mu^{t'}) = z) \\ &\leq \mathbb{P}(\cup_{t' \in \mathcal{T} \setminus \{t\}} \{m^t(X) - m^{t'}(X) + \mu^{t'} = z\}) \\ &\leq \sum_{t' \in \mathcal{T} \setminus \{t\}} \mathbb{P}(m^t(X) - m^{t'}(X) + \mu^{t'} = z) = 0, \end{aligned}$$

where the first inequality holds because the probability of  $\min_{t' \in \mathcal{T} \setminus \{t\}} (m^t(X) - m^{t'}(X) + \mu^{t'}) = z$  is less than or equal to the probability that  $m^t(X) - m^{t'}(X) + \mu^{t'} = z$  for at least one  $t' \in \mathcal{T}$ , the second inequality follows from the union bound, and the last equality holds because  $m^t(X) - m^{t'}(X)$  is continuously distributed for all  $t, t' \in \mathcal{T}$  by Assumption 2 (ii).  $\square$

In the remainder, we denote by  $\mathbb{P}_x$  the marginal distribution of  $X$  and by  $\mathbb{P}_{z^{\mu^{-t}}}$  the marginal distribution of  $Z^{\mu^{-t}}$  defined in Lemma EC.1 when necessary for clarity. The next lemma shows that the objective function of the dual problem  $(\mathcal{D})$  is differentiable under Assumption 2, and this result will later allow us to establish the optimality conditions for problem  $(\mathcal{D})$ .

LEMMA EC.2. *Define the auxiliary function  $f: \mathbb{R}^{m+1} \rightarrow \mathbb{R}$  through*

$$f(\mu) = \mathbb{E} \left[ \max_{t' \in \mathcal{T}} (m^{t'}(X) - \mu^{t'}) \right] = \int_{\mathcal{X}} \max_{t' \in \mathcal{T}} (m^{t'}(x) - \mu^{t'}) \, d\mathbb{P}_x(x), \quad (\text{EC.4})$$

*which appears in the objective function of the dual problem  $(\mathcal{D})$ . Under Assumption 2, the partial derivative of function  $f$  with respect to  $\mu^t$ ,  $t \in \mathcal{T}$ , is given by*

$$\frac{\partial}{\partial \mu^t} f(\mu) = -\mathbb{P}_x \left( m^t(X) - \mu^t \geq \max_{t' \in \mathcal{T} \setminus \{t\}} (m^{t'}(X) - \mu^{t'}) \right).$$

*Furthermore, each partial derivative is continuous in  $\mu$ , and  $f$  is therefore differentiable.*

*Proof.* We first derive the partial derivative of  $f$  with respect to each  $\mu^t$  and then show that each partial derivative is continuous in  $\mu \in \mathbb{R}_+^{m+1}$ . Given a real number  $a$ , we use  $a^+$  to denote the

positive part function  $\max(a, 0)$ . For any  $t \in \mathcal{T}$ , the partial derivative of  $f$  with respect to  $\mu^t$  is given by

$$\begin{aligned}
\frac{\partial}{\partial \mu^t} f(\mu) &= \frac{\partial}{\partial \mu^t} \int_{\mathcal{X}} \max_{t' \in \mathcal{T}} \left( m^{t'}(x) - \mu^{t'} \right) d\mathbb{P}_x(x) \\
&= \frac{\partial}{\partial \mu^t} \int_{\mathcal{X}} \max \left\{ m^t(x) - \mu^t, \max_{t' \in \mathcal{T} \setminus \{t\}} \left( m^{t'}(x) - \mu^{t'} \right) \right\} d\mathbb{P}_x(x) \\
&= \frac{\partial}{\partial \mu^t} \int_{\mathcal{X}} \left( m^t(x) - \max_{t' \in \mathcal{T} \setminus \{t\}} \left( m^{t'}(x) - \mu^{t'} \right) - \mu^t \right)^+ + \max_{t' \in \mathcal{T} \setminus \{t\}} \left( m^{t'}(x) - \mu^{t'} \right) d\mathbb{P}_x(x) \\
&= \frac{\partial}{\partial \mu^t} \int_{-\infty}^{\infty} (z - \mu^t)^+ d\mathbb{P}_{z^{\mu-t}}(z) \\
&= \frac{\partial}{\partial \mu^t} \int_{\mu^t}^{\infty} (z - \mu^t) d\mathbb{P}_{z^{\mu-t}}(z) \\
&= \int_{\mu^t}^{\infty} \frac{\partial}{\partial \mu^t} (z - \mu^t) d\mathbb{P}_{z^{\mu-t}}(z),
\end{aligned}$$

where the fourth equality follows from a variable substitution by using the definition of  $Z^{\mu-t}$  from Lemma EC.1 and that  $\frac{\partial}{\partial \mu^t} \int_{\mathcal{X}} \max_{t' \in \mathcal{T} \setminus \{t\}} \left( m^{t'}(x) - \mu^{t'} \right) d\mathbb{P}_x(x) = 0$ , and the last equality follows from the Leibniz integral rule, which applies because for all  $z \in \mathbb{R}$ ,  $\frac{\partial}{\partial \mu^t} (z - \mu^t) = -1$  exists and is continuous for all  $\mu^t \in \mathbb{R}_+$ , and it is also uniformly bounded by the constant 1. We thus have

$$\begin{aligned}
\frac{\partial}{\partial \mu^t} f(\mu) &= \int_{\mu^t}^{\infty} \frac{\partial}{\partial \mu^t} (z - \mu^t) d\mathbb{P}_{z^{\mu-t}}(z) \\
&= -\mathbb{P}_{z^{\mu-t}}(Z^{\mu-t} \geq \mu^t) \\
&= -\mathbb{P}_x \left( m^t(X) - \mu^t \geq \max_{t' \in \mathcal{T} \setminus \{t\}} \left( m^{t'}(X) - \mu^{t'} \right) \right),
\end{aligned} \tag{EC.5}$$

where the last equality follows again from the definition of  $Z^{\mu-t}$ .

Next, we show that  $\frac{\partial}{\partial \mu^t} f(\mu)$  is continuous in  $\mu$ , i.e.,  $\lim_{\mu \rightarrow \mu_0} \frac{\partial}{\partial \mu^t} f(\mu) = \frac{\partial}{\partial \mu^t} f(\mu_0)$  for every  $\mu_0 \in \mathbb{R}_+^{m+1}$ . To this end, consider an arbitrary  $\mu_0 \in \mathbb{R}_+^{m+1}$ , and note that we have

$$\lim_{\mu \rightarrow \mu_0} \frac{\partial}{\partial \mu^t} f(\mu) = \lim_{\mu \rightarrow \mu_0} - \int_{\mathcal{X}} \mathbb{1}[h^t(x, \mu) > 0] d\mathbb{P}_x(x),$$

where  $h^t(x, \mu) = m^t(x) - \max_{t' \in \mathcal{T} \setminus \{t\}} \left( m^{t'}(X) - \mu^{t'} \right) - \mu^t = Z^{\mu-t} - \mu^t$ , and the equality follows from (EC.5). The indicator function  $\mathbb{1}[h^t(x, \mu) > 0]$  is uniformly bounded by the constant 1. By Lemma EC.1,  $h^t(x, \mu)$  is continuously distributed for every  $\mu \in \mathbb{R}_+^{m+1}$  since  $Z^{\mu-t}$  is continuously distributed for any  $\mu^{-t} \in \mathbb{R}_+^m$ , and therefore  $\mathbb{P}_x(h^t(x, \mu) = 0) = 0$ . As the function  $h^t(x, \mu)$  is continuous in  $\mu$  for every  $x \in \mathcal{X}$ , the indicator function  $\mathbb{1}[h^t(x, \mu) > 0]$  is discontinuous at  $\mu_0$  only

if  $h(x, \mu_0) = 0$ . As  $\mathbb{P}_x(h(X, \mu_0) = 0) = 0$ , we have  $\lim_{\mu \rightarrow \mu_0} \mathbb{1}[h^t(X, \mu) > 0] = \mathbb{1}[h(X, \mu_0) > 0]$   $\mathbb{P}_x$ -almost surely. We thus obtain

$$\begin{aligned} \lim_{\mu \rightarrow \mu_0} \frac{\partial}{\partial \mu^t} f(\mu) &= \lim_{\mu \rightarrow \mu_0} - \int_{\mathcal{X}} \mathbb{1}[h^t(x, \mu) > 0] d\mathbb{P}_x(x) \\ &= - \int_{\mathcal{X}} \lim_{\mu \rightarrow \mu_0} \mathbb{1}[h^t(x, \mu) > 0] d\mathbb{P}_x(x) \\ &= - \int_{\mathcal{X}} \mathbb{1}[h(x, \mu_0) > 0] d\mathbb{P}_x(x) \\ &= \frac{\partial}{\partial \mu^t} f(\mu_0), \end{aligned}$$

where the second equality follows from the Dominated Convergence Theorem. This proves that the partial derivative  $\frac{\partial}{\partial \mu^t} f(\mu)$  is continuous in  $\mu$ . Function  $f$  is thus differentiable.  $\square$

Lemma EC.2 implies that the objective function of problem  $(\mathcal{D})$  is differentiable. As problem  $(\mathcal{D})$  is a convex optimization problem and satisfies Slater's condition, strong duality holds, i.e., the optimal value of  $(\mathcal{D})$  coincides with the one of its dual problem. Denote by  $\lambda$  the dual variable of the constraint  $\mu \geq 0$  in problem  $(\mathcal{D})$ . If  $\mu^*$  and  $\lambda^*$  are solutions to problem  $(\mathcal{D})$  and its dual, respectively, then  $\mu^*, \lambda^*$  satisfy the Karush–Kuhn–Tucker (KKT) conditions

$$\begin{aligned} \mathbb{P} \left( m^t(X) - \mu^{*,t} \geq \max_{t' \in \mathcal{T} \setminus \{t\}} (m^{t'}(X) - \mu^{*,t'}) \right) + \lambda^{*,t} &= b^t && \text{(Stationarity Condition),} \\ \lambda^{*,t} \mu^{*,t} &= 0 && \text{(Complementary Slackness),} \\ \mu^{*,t} \geq 0, \quad \lambda^{*,t} &\geq 0 && \text{(Primal \& Dual Feasibility),} \end{aligned} \tag{EC.6}$$

for all  $t \in \mathcal{T}$ . Intuitively, the KKT conditions imply that policy  $\pi^*$  defined in (EC.3) satisfies the capacity constraints in expectation. To see this, note that  $\mathbb{P} \left( m^t(X) - \mu^{*,t} \geq \max_{t' \in \mathcal{T} \setminus \{t\}} (m^{t'}(X) - \mu^{*,t'}) \right)$  is the expected number of individuals assigned treatment  $t$  under policy  $\pi^*$ . As  $\lambda^{*,t} \geq 0$  by dual feasibility, this quantity cannot exceed the available capacity  $b_t$  by the stationarity condition.

We are now ready to prove Theorem EC.1.

*Proof of Theorem EC.1.* For any  $\mu^*$  optimal in  $(\mathcal{D})$ , the treatment assignment probabilities of  $\pi^*$  defined in (EC.3) and that of the dual-price queuing policy defined in Definition EC.1 coincide. Indeed, the dual-price queuing policy is nothing more than an online implementation of  $\pi^*$  since both make assignments using the same treatment assignment criteria. The difference is that the dual-price queuing policy waitlists individuals for treatments by assigning them to queues since there may be a mismatch in timing between the arrival of individuals and resources. It is thus sufficient to show that  $\pi^*$  is optimal in  $(\mathcal{P})$ . We first prove that  $\pi^*$  is feasible in  $(\mathcal{P})$  and then show that its expected average outcome matches the optimal value  $z^*$  of  $(\mathcal{P})$ .

Policy  $\pi^*$  is feasible in  $(\mathcal{P})$  because

$$\begin{aligned}\mathbb{E}[\pi^{*,t}(X)] &= \mathbb{P}\left(t = \min \operatorname{argmax}_{t' \in \mathcal{T}} m^{t'}(X) - \mu^{*,t'}\right) \\ &\leq \mathbb{P}\left(m^t(X) - \mu^{*,t} \geq \max_{t' \in \mathcal{T} \setminus \{t\}} (m^{t'}(X) - \mu^{*,t'})\right) \leq b^t \quad \forall t \in \mathcal{T},\end{aligned}$$

where the first inequality follows from the fact that if the event  $t = \min \operatorname{argmax}_{t' \in \mathcal{T}} m^{t'}(X) - \mu^{*,t'}$  occurs, then the inequality  $m^t(X) - \mu^{*,t} \geq \max_{t' \in \mathcal{T} \setminus \{t\}} (m^{t'}(X) - \mu^{*,t'})$  must hold, and the second inequality follows from the stationarity and dual feasibility conditions in (EC.6).

As  $\pi^*$  is feasible in  $(\mathcal{P})$ , its expected average outcome cannot exceed  $z^*$ . We next show that the expected average outcome of  $\pi^*$  is at least as high as  $z^*$ , which implies that the two values coincide. The expected average outcome of  $\pi^*$  amounts to

$$\begin{aligned}\mathbb{E}\left[\sum_{t \in \mathcal{T}} \mathbb{1}\left[t = \min \operatorname{argmax}_{t' \in \mathcal{T}} m^{t'}(X) - \mu^{*,t'}\right] m^t(X)\right] \\ &= \mathbb{E}\left[\max_{t \in \mathcal{T}} (m^t(X) - \mu^{*,t}) + \sum_{t \in \mathcal{T}} \mathbb{1}\left[t = \min \operatorname{argmax}_{t' \in \mathcal{T}} m^{t'}(X) - \mu^{*,t'}\right] \mu^{*,t}\right] \\ &= \mathbb{E}\left[\max_{t \in \mathcal{T}} (m^t(X) - \mu^{*,t})\right] + \sum_{t \in \mathcal{T}} \mathbb{P}\left(t = \min \operatorname{argmax}_{t' \in \mathcal{T}} m^{t'}(X) - \mu^{*,t'}\right) \mu^{*,t} \\ &= \mathbb{E}\left[\max_{t \in \mathcal{T}} (m^t(X) - \mu^{*,t})\right] + \sum_{t \in \mathcal{T}} b^t \mu^{*,t} \\ &= \nu^* \geq z^*.\end{aligned}\tag{EC.7}$$

The third equality above follows from the KKT conditions (EC.6), which imply that

$$\begin{aligned}\mu^{*,t} b^t &= \mu^{*,t} \left[ \mathbb{P}\left(m^t(X) - \mu^{*,t} \geq \max_{t' \in \mathcal{T} \setminus \{t\}} (m^{t'}(X) - \mu^{*,t'})\right) + \lambda^{*,t} \right] \\ &= \mu^{*,t} \mathbb{P}\left(m^t(X) - \mu^{*,t} \geq \max_{t' \in \mathcal{T} \setminus \{t\}} (m^{t'}(X) - \mu^{*,t'})\right) \\ &= \mu^{*,t} \mathbb{P}\left(t = \min \operatorname{argmax}_{t' \in \mathcal{T}} m^{t'}(X) - \mu^{*,t'}\right),\end{aligned}$$

where the first and second equalities follows from the stationarity and complementary slackness conditions, respectively, and the third equality holds because  $\operatorname{argmax}_{t' \in \mathcal{T}} m^{t'}(X) - \mu^{*,t}$  is a singleton almost surely. The last equality in (EC.7) holds because  $\mu^*$  is optimal in  $(\mathcal{D})$ , and the inequality follows from weak duality. The claim thus follows.  $\square$

### EC.2.2. Proofs for Section 3

The proof of Theorem 1 requires the following technical lemmas: Lemmas EC.3 – EC.7.

We first define the following functions  $\nu : \mathbb{R}^{m+1} \rightarrow \mathbb{R}$  and  $\hat{\nu}_n : \mathbb{R}^{m+1} \rightarrow \mathbb{R}$  for denoting the objective functions of the dual problem  $(\mathcal{D})$  and its sample approximation  $(\hat{\mathcal{D}})$  as functions of  $\mu$ . These definitions will be used throughout Section EC.2.2.

$$\begin{aligned}\nu(\mu) &= \mathbb{E} \left[ \max_{t \in \mathcal{T}} (m^t(X) - \mu^t) \right] + \sum_{t \in \mathcal{T}} \mu^t b^t \\ \hat{\nu}_n(\mu) &= \frac{1}{n} \sum_{i=1}^n \max_{t \in \mathcal{T}} (\hat{m}_n^t(x_i) - \mu^t) + \sum_{t \in \mathcal{T}} \mu^t b^t.\end{aligned}\tag{EC.8}$$

By definition, we have  $\nu^* = \min_{\mu \in \mathbb{R}_+^{m+1}} \nu(\mu)$  and  $\hat{\nu}_n^* = \min_{\mu \in \mathbb{R}_+^{m+1}} \hat{\nu}_n(\mu)$ . Denote by  $\mathcal{S}^*$  and  $\hat{\mathcal{S}}_n^*$  the sets of optimal solutions to problems  $(\mathcal{D})$  and  $(\hat{\mathcal{D}})$ , respectively. The next lemma shows that  $\mathcal{S}^*$  is contained in a compact set as long as the function  $m^t$  is uniformly bounded for all  $t \in \mathcal{T}$ , which is guaranteed under Assumption 2.

**LEMMA EC.3.** *Under Assumption 2, the set  $\mathcal{S}^*$  of optimal solutions to problem  $(\mathcal{D})$  is non-empty and contained in the compact set  $[0, C]^{m+1} \subset \mathbb{R}_+^{m+1}$ , where  $C$  is defined as in Assumption 2(i), i.e.,  $|m^t(x)| \leq C$  for all  $x \in \mathcal{X}$  and  $t \in \mathcal{T}$ .*

*Proof.* We will first show that the optimal solution set  $\mathcal{S}^*$  is contained in the compact set  $[0, C]^{m+1}$  and then show that the set  $\mathcal{S}^*$  is non-empty.

Suppose for contradiction that  $\mathcal{S}^* \not\subseteq [0, C]^{m+1}$ , i.e., there exists a  $\mu \in \mathcal{S}^*$  such that  $\mu \notin [0, C]^{m+1}$ . As  $\mu \notin [0, C]^{m+1}$ , there exists at least one  $t \in \mathcal{T}$  such that  $\mu^t > C$ . Denote by  $\mathcal{T}' = \{t \in \mathcal{T} : \mu^t > C\}$  the nonempty set of treatments  $t$  for which  $\mu^t > C$ . We construct another feasible solution  $\mu' \in [0, C]^{m+1}$  via  $(\mu')^t = \mu^t$  for all  $t \in \mathcal{T} \setminus \mathcal{T}'$  (all  $t$ 's where  $\mu^t$  is in the interval  $[0, C]$ ), and  $(\mu')^t = C$  for all  $t \in \mathcal{T}'$  (all  $t$ 's where  $\mu^t$  is not in the interval  $[0, C]$ ). We will show that the objective value of  $\mu'$  in problem  $(\mathcal{D})$  is strictly lower than that of  $\mu$ . This implies that  $\mu$  cannot be optimal and results in a contradiction.

We will first show that  $\max_{t \in \mathcal{T}} (m^t(x) - \mu^t) = \max_{t \in \mathcal{T}} (m^t(x) - (\mu')^t)$  for all  $x \in \mathcal{X}$ . To this end, recall that treatment 0 represents the no-treatment option, and  $b^0 = 1$ , which makes the capacity constraint corresponding to this treatment redundant. This implies that we can assume without loss of generality  $\mu^0 = 0$ , and therefore  $0 \notin \mathcal{T}'$ . For any  $t \in \mathcal{T}'$  and  $x \in \mathcal{X}$ , we have

$$m^t(x) - \mu^t < m^t(x) - C = m^t(x) - (\mu')^t \leq m^0(x) = m^0(x) - \mu^0 = m^0(x) - (\mu')^0, \tag{EC.9}$$

where the first inequality follows from the definition of  $\mathcal{T}'$ , the first equality follows from the definition of  $\mu'$ , and the second inequality holds because  $m^t(x) \leq C$  and  $m^0(x)$  is non-negative. The last two equalities hold because  $\mu^0 = (\mu')^0 = 0$ . By equation (EC.9), we have  $m^t(x) - \mu^t < m^0(x) - \mu^0$  and  $m^t(x) - (\mu')^t \leq m^0(x) - (\mu')^0$  for all  $t \in \mathcal{T}'$ . This implies that for any  $x \in \mathcal{X}$ , we have

$$\max_{t \in \mathcal{T}} (m^t(x) - \mu^t) = \max_{t \in \mathcal{T} \setminus \mathcal{T}'} (m^t(x) - \mu^t) = \max_{t \in \mathcal{T} \setminus \mathcal{T}'} (m^t(x) - (\mu')^t) = \max_{t \in \mathcal{T}} (m^t(x) - (\mu')^t),$$

where the first and third equality hold because no  $t \in \mathcal{T}'$  can attain the maximum in  $\max_{t \in \mathcal{T}}(m^t(x) - \mu^t)$  and  $\max_{t \in \mathcal{T}}(m^t(x) - (\mu')^t)$  by (EC.9), and the second equality holds because  $(\mu')^t = \mu^t$  for all  $t \in \mathcal{T} \setminus \mathcal{T}'$ .

The objective value  $\nu(\mu)$  of  $\mu$  thus exceeds the objective value  $\nu(\mu')$  of  $\mu'$  as

$$\begin{aligned} \nu(\mu) &= \mathbb{E} \left[ \max_{t \in \mathcal{T}}(m^t(X) - \mu^t) \right] + \sum_{t \in \mathcal{T}} \mu^t b^t = \mathbb{E} \left[ \max_{t \in \mathcal{T}}(m^t(X) - (\mu')^t) \right] + \sum_{t \in \mathcal{T}} \mu^t b^t \\ &> \mathbb{E} \left[ \max_{t \in \mathcal{T}}(m^t(X) - (\mu')^t) \right] + \sum_{t \in \mathcal{T}} (\mu')^t b^t \\ &= \nu(\mu'), \end{aligned}$$

where the second equality holds because we showed that  $\max_{t \in \mathcal{T}}(m^t(x) - \mu^t) = \max_{t \in \mathcal{T}}(m^t(x) - (\mu')^t)$  for all  $x \in \mathcal{X}$ , and the inequality holds because  $(\mu')^t \leq \mu^t$  for all  $t \in \mathcal{T}$  and the inequality is strict for all  $t \in \mathcal{T}'$  by definition of  $\mu'$ . This implies that  $\mu$  cannot be optimal and results in a contradiction. Thus, we must have  $\mathcal{S}^* \subseteq [0, C]^{m+1}$ .

We now show that  $\mathcal{S}^*$  is always non-empty. As  $\mathcal{S}^* \subseteq [0, C]^{m+1}$ , the dual problem  $(\mathcal{D})$  is equivalent to  $\min_{\mu \in [0, C]^{m+1}} \nu(\mu)$ , where we only minimize over the compact set  $[0, C]^{m+1}$ . Since  $\nu(\mu)$  is a convex function, and hence continuous, in  $\mathbb{R}^n$ , it attains a minimum over any compact set (Proposition A.2.7 of Bertsekas (2009)), which implies  $\mathcal{S}^*$  is non-empty.  $\square$

Similar to Lemma EC.3, the next lemma shows that the optimal solution set  $\hat{\mathcal{S}}_n^*$  must also lie within a compact set.

**LEMMA EC.4.** *Under Assumption 3, for any  $n \in \mathbb{N}$ , the set  $\hat{\mathcal{S}}_n^*$  is non-empty and contained in the compact set  $[0, \hat{C}]^{m+1} \subset \mathbb{R}_+^{m+1}$ , where  $\hat{C}$  is defined as in Assumption 3(i), i.e.,  $|\hat{m}_n^t(x)| \leq \hat{C}$  for all  $x \in \mathcal{X}$ ,  $t \in \mathcal{T}$ , and  $n \in \mathbb{N}$ .*

*Proof.* Fix an arbitrary  $\omega \in \Omega$ , i.e., fix historical samples. By Assumption 3(i),  $\hat{C}$  is a uniform upper bound on  $|\hat{m}_n^t(x)|$  for all  $x \in \mathcal{X}$ ,  $t \in \mathcal{T}$ , and  $n \in \mathbb{N}$ . Following the same arguments used in the proof of Lemma EC.3, we can show that any solution  $\mu \notin [0, \hat{C}]^{m+1}$  will be sub-optimal and there exists a solution in  $[0, \hat{C}]^{m+1}$  with strictly lower objective value. Therefore, the feasible set of problem  $(\hat{\mathcal{D}})$  can be restricted to the compact set  $[0, \hat{C}]^{m+1}$ . For all  $n \in \mathbb{N}$ ,  $\hat{\nu}_n(\mu)$  is a convex function and hence attains a minimum over any compact set. Since the choice of  $\omega$  was arbitrary, for any  $\omega \in \Omega$ ,  $\hat{\mathcal{S}}_n^*$  is non-empty for all  $n \in \mathbb{N}$  and contained in  $[0, \hat{C}]^{m+1}$ .  $\square$

The next lemma shows that the objective function  $\hat{\nu}_n(\mu)$  of the sample-approximate dual problem  $(\hat{\mathcal{D}})$  converges almost surely to the objective function  $\nu(\mu)$  of the true dual problem  $(\mathcal{D})$  uniformly in  $\mu$ .

**LEMMA EC.5.** *Under Assumptions 2 and 3, the objective function  $\hat{\nu}_n(\mu)$  of problem  $(\hat{\mathcal{D}})$  converges almost surely to the objective function  $\nu(\mu)$  of problem  $(\mathcal{D})$  uniformly on any nonempty compact set  $\mathcal{S} \subseteq \mathbb{R}_+^{m+1}$  as  $n$  tends to infinity, that is,*

$$\lim_{n \rightarrow \infty} \sup_{\mu \in \mathcal{S}} |\nu(\mu) - \hat{\nu}_n(\mu)| = 0.$$

*Proof.* Consider an arbitrary nonempty compact set  $\mathcal{S} \subseteq \mathbb{R}_+^{m+1}$ . By triangle inequality, we have

$$\begin{aligned} \sup_{\mu \in \mathcal{S}} |\nu(\mu) - \hat{\nu}_n(\mu)| &\leq \sup_{\mu \in \mathcal{S}} \left( |\nu(\mu) - \nu'_n(\mu)| + |\nu'_n(\mu) - \hat{\nu}_n(\mu)| \right) \\ &\leq \sup_{\mu \in \mathcal{S}} |\nu(\mu) - \nu'_n(\mu)| + \sup_{\mu \in \mathcal{S}} |\nu'_n(\mu) - \hat{\nu}_n(\mu)|, \end{aligned}$$

where  $\nu'_n$  is a sample average approximation of  $\nu$  using the true  $m^t$  functions, i.e.,

$$\nu'_n(\mu) = \frac{1}{n} \sum_{i=1}^n \max_{t \in \mathcal{T}} (m^t(x_i) - \mu^t) + \sum_{t \in \mathcal{T}} \mu^t b^t.$$

In the following, we investigate the terms  $\sup_{\mu \in \mathcal{S}} |\nu(\mu) - \nu'_n(\mu)|$  and  $\sup_{\mu \in \mathcal{S}} |\nu'_n(\mu) - \hat{\nu}_n(\mu)|$  one by one.

We first consider the term  $\sup_{\mu \in \mathcal{S}} |\nu(\mu) - \nu'_n(\mu)|$  and show that

$$\lim_{n \rightarrow \infty} \sup_{\mu \in \mathcal{S}} |\nu(\mu) - \nu'_n(\mu)| = 0. \quad (\text{EC.10})$$

To this end, let  $F(x, \mu) = \max_{t \in \mathcal{T}} (m^t(x) - \mu^t)$ , and note that, for any  $x \in \mathcal{X}$ ,  $F(x, \mu)$  is continuous in  $\mu$  as it is a piecewise-linear function of  $\mu$ . For any  $\mu \in \mathcal{S}$ , the function  $F(x, \mu)$  is dominated by the integrable function  $g(x) = C$  as  $C$  is a uniform bound on  $m^t$  for each  $t \in \mathcal{T}$  by Assumption 2(i). We can thus invoke Theorem 7.53 of Shapiro et al. (2014) and conclude that  $\frac{1}{n} \sum_{i=1}^n F(X_i, \mu)$  converges uniformly on  $\mathcal{S}$  and almost surely to  $\mathbb{E}[F(X, \mu)]$  as  $n$  tends to infinity. This implies that  $\nu'_n(\mu)$  converges uniformly on  $\mathcal{S}$  and almost surely to  $\nu(\mu)$ . We thus proved (EC.10).

Next, we consider the term  $\sup_{\mu \in \mathcal{S}} |\nu'_n(\mu) - \hat{\nu}_n(\mu)|$  and similarly show that

$$\lim_{n \rightarrow \infty} \sup_{\mu \in \mathcal{S}} |\nu'_n(\mu) - \hat{\nu}_n(\mu)| = 0. \quad (\text{EC.11})$$

For any  $\mu \in \mathcal{S}$ , we have

$$\begin{aligned} |\nu'_n(\mu) - \hat{\nu}_n(\mu)| &= \left| \frac{1}{n} \sum_{i=1}^n \max_{t \in \mathcal{T}} (m^t(x_i) - \mu^t) - \max_{t \in \mathcal{T}} (\hat{m}_n^t(x_i) - \mu^t) \right| \\ &\leq \frac{1}{n} \sum_{i=1}^n \left| \max_{t \in \mathcal{T}} (m^t(x_i) - \mu^t) - \max_{t \in \mathcal{T}} (\hat{m}_n^t(x_i) - \mu^t) \right| \\ &\leq \frac{1}{n} \sum_{i=1}^n \left| \max_{t \in \mathcal{T}} (m^t(x_i) - \hat{m}_n^t(x_i)) \right| \\ &\leq \frac{1}{n} \sum_{i=1}^n \max_{t \in \mathcal{T}} |m^t(x_i) - \hat{m}_n^t(x_i)| \\ &\leq \sup_{x \in \mathcal{X}} \max_{t \in \mathcal{T}} |m^t(x) - \hat{m}_n^t(x)|. \end{aligned} \quad (\text{EC.12})$$

The first inequality follows by triangle inequality. To see that the second inequality holds, fix an arbitrary  $x \in \mathcal{X}$  and  $\mu \in \mathcal{S}$  and let  $t^* = \min \arg \max_{t \in \mathcal{T}} (m^t(x) - \mu^t)$ . We then have

$$\begin{aligned} \max_{t \in \mathcal{T}} (m^t(x) - \mu^t) - \max_{t \in \mathcal{T}} (\hat{m}_n^t(x) - \mu^t) &\leq (m^{t^*}(x) - \mu^{t^*}) - (\hat{m}_n^{t^*}(x) - \mu^{t^*}) \\ &= m^{t^*}(x) - \hat{m}_n^{t^*}(x) \\ &\leq \max_{t \in \mathcal{T}} [m^t(x) - \hat{m}_n^t(x)], \end{aligned}$$

where the first inequality holds by definition of  $t^*$  and  $\hat{m}_n^{t^*}(x) - \mu^{t^*} \leq \max_{t \in \mathcal{T}} (\hat{m}_n^t(x) - \mu^t)$ . Since the choice of  $x$  and  $\mu$  were arbitrary, the above holds for all  $x \in \mathcal{X}$  and  $\mu \in \mathcal{S}$ , thereby proving the second inequality in (EC.12). Finally, by Assumption 3(ii), we have  $\lim_{n \rightarrow \infty} \sup_{x \in \mathcal{X}} \max_{t \in \mathcal{T}} |m^t(x) - \hat{m}_n^t(x)| = 0$ . As this claim holds irrespective of the value of  $\mu \in \mathcal{S}$  and in view of (EC.12), we conclude that (EC.11) holds.

The findings above imply that  $\lim_{n \rightarrow \infty} \sup_{\mu \in \mathcal{S}} |\nu(\mu) - \hat{\nu}_n(\mu)| = 0$ . The claim thus follows.  $\square$

Next, we use Lemmas EC.3–EC.5 to show that the optimal value  $\hat{\nu}_n^*$  of the sample-approximate dual problem  $(\hat{\mathcal{D}})$  converges almost surely to the optimal value  $\nu^*$  of the true dual problem  $(\mathcal{D})$ . We will also prove a convergence result in terms of their respective optimal solution sets. The following definitions will be relevant for this result. We use  $\|\cdot\|$  to denote the Euclidean norm and define

$$\text{dist}(x, \mathcal{A}) = \inf_{x' \in \mathcal{A}} \|x - x'\|, \quad \mathbb{D}(\mathcal{A}, \mathcal{B}) = \sup_{x \in \mathcal{A}} \text{dist}(x, \mathcal{B}),$$

where  $\text{dist}(x, \mathcal{A})$  denotes the distance between a point  $x \in \mathbb{R}^d$  and a set  $\mathcal{A} \subseteq \mathbb{R}^d$ , and  $\mathbb{D}(\mathcal{A}, \mathcal{B})$  denotes the deviation between two sets  $\mathcal{A} \subseteq \mathbb{R}^d$  and  $\mathcal{B} \subseteq \mathbb{R}^d$ .

**LEMMA EC.6.** *Under Assumptions 2 and 3, the optimal value  $\hat{\nu}_n^*$  of problem  $(\hat{\mathcal{D}})$  converges almost surely to the optimal value  $\nu^*$  of problem  $(\mathcal{D})$ , and the deviation  $\mathbb{D}(\hat{\mathcal{S}}_n^*, \mathcal{S}^*)$  between their respective optimal solution sets  $\hat{\mathcal{S}}_n^*$  and  $\mathcal{S}^*$  converges almost surely to zero as  $n$  tends to infinity, that is,  $\lim_{n \rightarrow \infty} |\nu^* - \hat{\nu}_n^*| = 0$  and  $\lim_{n \rightarrow \infty} \mathbb{D}(\hat{\mathcal{S}}_n^*, \mathcal{S}^*) = 0$ .*

*Proof.* Consider the compact set  $[0, \max(C, \hat{C})]^{m+1} \subset \mathbb{R}_+^{m+1}$ , and note that: (i) the optimal solution set  $\mathcal{S}^*$  is nonempty and contained in  $[0, \max(C, \hat{C})]^{m+1}$  by Lemma EC.3; (ii) the function  $\nu(\mu)$  is finite valued by Assumption 2(i) and continuous on  $[0, \max(C, \hat{C})]^{m+1}$ ; (iii)  $\hat{\nu}_n(\mu)$  converges uniformly on  $[0, \max(C, \hat{C})]^{m+1}$  and almost surely to  $\nu(\mu)$  by Lemma EC.5; and, (iv)  $\hat{\mathcal{S}}_n^*$  is nonempty and contained in  $[0, \max(C, \hat{C})]^{m+1}$  by Lemma EC.4. We can thus invoke Theorem 5.3 of Shapiro et al. (2014) and conclude that  $\hat{\nu}_n^*$  converges almost surely to  $\nu^*$ , and  $\mathbb{D}(\hat{\mathcal{S}}_n^*, \mathcal{S}^*)$  converges almost surely to zero.  $\square$

The proof of Theorem 1 will require comparing the treatment assignments made by a dual-price queuing policy and its sample approximation defined in Definitions EC.1 and 1, respectively. To that end, we extend the sample solution set convergence results of Lemma EC.6 by showing that for  $n$  large enough, any optimal sample-based solution  $\hat{\mu}_n^*$  will be arbitrarily close to some element within the optimal solution set  $\mathcal{S}^*$ . This will help ensure that almost surely for  $n$  large enough, a sample approximation policy will make the same treatment assignments as a dual-price queuing policy.

**LEMMA EC.7.** *Under Assumptions 2 and 3, there exists  $\Omega' \subseteq \Omega$ , where  $\mathbb{P}(\Omega') = 1$ , with the following property. For any  $\omega \in \Omega'$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that*

$$\forall \hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega), \exists \mu^* \in \mathcal{S}^* : \|\hat{\mu}_n^*(\omega) - \mu^*\| < \epsilon \quad \forall n \geq N(\omega).$$



*Proof.* By Lemma EC.6,  $\mathbb{D}(\hat{\mathcal{S}}_n^*, \mathcal{S}^*)$  converges almost surely to zero as  $n \rightarrow \infty$ . In other words, there exists  $\Omega' \subseteq \Omega$ , where  $\mathbb{P}(\Omega') = 1$ , with the following property. For any  $\omega \in \Omega'$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that

$$\sup_{\hat{\mu}_n(\omega) \in \hat{\mathcal{S}}_n^*(\omega)} \inf_{\mu \in \mathcal{S}^*} \|\hat{\mu}_n(\omega) - \mu\| < \epsilon \quad \forall n \geq N(\omega),$$

which follows by the definition of  $\mathbb{D}(A, B)$ . The above implies that

$$\forall \hat{\mu}_n(\omega) \in \hat{\mathcal{S}}_n^*(\omega) : \inf_{\mu \in \mathcal{S}^*} \|\hat{\mu}_n(\omega) - \mu\| < \epsilon \quad \forall n \geq N(\omega). \quad (\text{EC.13})$$

Next, we show that for every  $\omega \in \Omega$  (i.e., for fixed historical sample),  $\inf_{\mu \in \mathcal{S}^*} \|\hat{\mu}_n(\omega) - \mu\|$  has an optimal solution for all  $\hat{\mu}_n(\omega) \in \hat{\mathcal{S}}_n^*(\omega)$ . To this end, we first show that  $\mathcal{S}^*$  is a compact set. Recall that  $\mathcal{S}^*$  denotes the optimal solution set to problem  $(\mathcal{D})$ . The set  $\mathcal{S}^*$  is bounded as it is contained in the compact set  $[0, C]^{m+1}$  by Lemma EC.3, where  $C$  is defined as in Assumption 2(i). The objective function  $\nu(\mu)$  of problem  $(\mathcal{D})$  is convex and thus continuous on  $\mathbb{R}^{m+1}$ . This implies that it has closed sublevel sets of the form  $\{\mu \in \mathbb{R}_+^{m+1} \mid \nu(\mu) \leq c\}$ , where  $c$  is any scalar. We now have  $\mathcal{S}^* = [0, C]^{m+1} \cap \{\mu \in \mathbb{R}_+^{m+1} \mid \nu(\mu) \leq \nu^*\}$ , which is closed as it is an intersection of closed sets. The set  $\mathcal{S}^*$  is thus compact.

Since all norms are continuous on  $\mathbb{R}_+^{m+1}$  and  $\mathcal{S}^*$  is compact, by Weierstrass' Extreme Value Theorem (Proposition A.2.7 of Bertsekas (2009)),  $\inf_{\mu \in \mathcal{S}^*} \|\hat{\mu}_n(\omega) - \mu\|$  has an optimal solution for any input  $\hat{\mu}_n(\omega) \in \hat{\mathcal{S}}_n^*(\omega)$ . This together with (EC.13) implies that for any  $\omega \in \Omega'$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that

$$\forall \hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega), \exists \mu^* \in \mathcal{S}^* : \|\hat{\mu}_n^* - \mu^*\| < \epsilon \quad \forall n \geq N(\omega).$$

The claim thus follows as  $\mathbb{P}(\Omega') = 1$ .  $\square$

We are now ready to prove Theorem 1.

*Proof of Theorem 1.* We use the following notation throughout the proof. Define the function  $h^t$  through  $h^t(x, \mu) = m^t(x) - \mu^t$  for all  $x \in \mathcal{X}$ ,  $\mu \in \mathbb{R}_+^{m+1}$ , and  $t \in \mathcal{T}$ . For any  $\omega \in \Omega$  (i.e., fixed historical samples), define the function  $\hat{h}_n^t(\cdot, \omega)$  through  $\hat{h}_n^t(x, \mu, \omega) = \hat{m}_n^t(x, \omega) - \mu^t$  for all  $x \in \mathcal{X}$ ,  $\mu \in \mathbb{R}_+^{m+1}$ ,  $t \in \mathcal{T}$ , and  $n \in \mathbb{N}$ . Denote by  $z^D$  the expected average outcome of a dual-price queuing policy. Recall that  $\hat{z}_n^D$  and  $z^*$  denote the expected average outcome of a sample-based dual-price queuing policy, where the expectation is taken with respect to the distribution of covariates and outcomes of the individuals arriving during implementation and *not* with respect to the distribution of the historical samples, and the optimal value of problem  $(\mathcal{P})$ , respectively. Note that, by definition,  $\hat{z}_n^D$  is a random variable as it depends on the historical samples observed.

The proof is divided into two steps. In Step 1, we show that the allocations of a sample-based dual-price queuing policy match almost surely those of a dual-price queuing policy as  $n \rightarrow \infty$ . In

Step 2, we then use this fact to show that a sample-based dual-price queuing policy asymptotically generates at least as high expected average outcome as that of a dual-price queuing policy as  $n \rightarrow \infty$  almost surely. As the expected average outcome of any dual-price queuing policy is equal to  $z^*$  by Theorem EC.1, we will thus conclude that the claim holds.

**Step 1.** For any  $\omega \in \Omega$ ,  $n \in \mathbb{N}$ ,  $\hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega)$ , and  $\mu^* \in \mathcal{S}^*$ , by triangle inequality we have

$$\begin{aligned} \sup_{t \in \mathcal{T}} \sup_{x \in \mathcal{X}} |\hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) - h^t(x, \mu^*)| &= \sup_{t \in \mathcal{T}} \sup_{x \in \mathcal{X}} |\hat{m}_n^t(x, \omega) - \hat{\mu}_n^{*,t}(\omega) - m^t(x) + \mu^{*,t}| \\ &\leq \sup_{t \in \mathcal{T}} \sup_{x \in \mathcal{X}} |\hat{m}_n^t(x, \omega) - m^t(x)| + \sup_{t \in \mathcal{T}} |\hat{\mu}_n^{*,t}(\omega) - \mu^{*,t}|. \end{aligned} \quad (\text{EC.14})$$

For the first term in the last line of (EC.14), by Assumption 3(ii), there exists  $\Omega' \subseteq \Omega$ , where  $\mathbb{P}(\Omega') = 1$ , with the following property. For every  $\omega \in \Omega'$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that

$$\sup_{t \in \mathcal{T}} \sup_{x \in \mathcal{X}} |\hat{m}_n^t(x, \omega) - m^t(x)| < \epsilon \quad \forall n \geq N(\omega).$$

For the second term in the last line of (EC.14), Lemma EC.7 and the property  $\sup_{t \in \mathcal{T}} |\mu^t| \leq \|\mu\|$  imply that there exists  $\Omega'' \subseteq \Omega$ , where  $\mathbb{P}(\Omega'') = 1$ , with the following property. For every  $\omega \in \Omega''$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that

$$\forall \hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega), \exists \mu^* \in \mathcal{S}^* : \sup_{t \in \mathcal{T}} |\hat{\mu}_n^{*,t}(\omega) - \mu^{*,t}| < \epsilon \quad \forall n \geq N(\omega),$$

Let  $\hat{\Omega} = \Omega' \cap \Omega''$  and note that  $\mathbb{P}(\hat{\Omega}) = 1$  because

$$\mathbb{P}(\hat{\Omega}) \geq 1 - \mathbb{P}((\Omega \setminus \Omega') \cup (\Omega \setminus \Omega'')) \geq 1 - \mathbb{P}((\Omega \setminus \Omega')) - \mathbb{P}((\Omega \setminus \Omega'')) = 1,$$

where the second inequality follows from the union bound, and the equality holds because  $\mathbb{P}(\Omega') = 1$  and  $\mathbb{P}(\Omega'') = 1$ . Assumption 3(ii) and Lemma EC.7 thus imply that for every  $\omega \in \hat{\Omega}$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that

$$\forall \hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega), \exists \mu^* \in \mathcal{S}^* : \sup_{t \in \mathcal{T}} \sup_{x \in \mathcal{X}} |\hat{m}_n^t(x, \omega) - m^t(x)| + \sup_{t \in \mathcal{T}} |\hat{\mu}_n^{*,t}(\omega) - \mu^{*,t}| < \epsilon \quad \forall n \geq N(\omega).$$

The above expression, in conjunction with equation (EC.14), imply that for every  $\omega \in \hat{\Omega}$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that

$$\forall \hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega), \exists \mu^* \in \mathcal{S}^* : \sup_{t \in \mathcal{T}} \sup_{x \in \mathcal{X}} |\hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) - h^t(x, \mu^*)| < \epsilon \quad \forall n \geq N(\omega). \quad (\text{EC.15})$$

Decomposing the expression  $\sup_{t \in \mathcal{T}} \sup_{x \in \mathcal{X}} |\hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) - h^t(x, \mu^*)| < \epsilon$  in Equation (EC.15), we have that for every  $\omega \in \hat{\Omega}$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that

$$\begin{aligned} \forall \hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega), \exists \mu^* \in \mathcal{S}^* : \hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) - h^t(x, \mu^*) < \epsilon, \quad h^t(x, \mu^*) - \hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) < \epsilon \\ \forall t \in \mathcal{T}, \forall x \in \mathcal{X}, \forall n \geq N(\omega). \end{aligned} \quad (\text{EC.16})$$

For any  $t, t' \in \mathcal{T}$ , by summing the inequalities from above, i.e.,

$$h^t(x, \mu^*) - \hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) < \epsilon \quad \text{and} \quad \hat{h}_n^{t'}(x, \hat{\mu}_n^*(\omega), \omega) - h^{t'}(x, \mu^*) < \epsilon,$$

we obtain

$$h^t(x, \mu^*) - \hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) + \hat{h}_n^{t'}(x, \hat{\mu}_n^*(\omega), \omega) - h^{t'}(x, \mu^*) < 2\epsilon. \quad (\text{EC.17})$$

Using (EC.16) and (EC.17), we have that for every  $\omega \in \hat{\Omega}$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that

$$\begin{aligned} \forall \hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega), \exists \mu^* \in \mathcal{S}^* : h^t(x, \mu^*) - \hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) + \hat{h}_n^{t'}(x, \hat{\mu}_n^*(\omega), \omega) - h^{t'}(x, \mu^*) < 2\epsilon \\ \forall t, t' \in \mathcal{T}, \forall x \in \mathcal{X}, \forall n \geq N(\omega). \end{aligned} \quad (\text{EC.18})$$

Let  $t(x, \mu) = \min \arg \max_{t \in \mathcal{T}} h^t(x, \mu)$ , and define  $\delta(x, \mu) = h^{t(x, \mu)}(x, \mu) - \max_{t \neq t(x, \mu)} h^t(x, \mu)$  for all  $x \in \mathcal{X}$  and  $\mu \in \mathbb{R}_+^{m+1}$ . Note that  $\delta(x, \mu)$  represents the difference between the highest and second highest value of the set  $\{h^t(x, \mu)\}_{t \in \mathcal{T}}$  and that we have  $\delta(x, \mu) > 0$  if and only if  $\arg \max_{t \in \mathcal{T}} h^t(x, \mu) = \{t(x, \mu)\}$ , i.e.,  $t(x, \mu)$  is the unique treatment that attains the maximum. Next, we write (EC.18) not for all treatment pairs  $t, t' \in \mathcal{T}$  but only for the treatment pairs of the form  $(t(x, \mu^*), t)$ , where  $t \in \mathcal{T} \setminus \{t(x, \mu^*)\}$ . Using the new notation we introduced and (EC.18), we have that for every  $\omega \in \hat{\Omega}$  and  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$

$$\begin{aligned} \forall \hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega), \exists \mu^* \in \mathcal{S}^* : \hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) - \hat{h}_n^{t(x, \mu^*)}(x, \hat{\mu}_n^*(\omega), \omega) < 2\epsilon - \delta(x, \mu^*) \\ \forall t \in \mathcal{T} \setminus \{t(x, \mu^*)\}, \forall x \in \mathcal{X}, \forall n \geq N(\omega), \end{aligned} \quad (\text{EC.19})$$

where we used the fact that

$$\delta(x, \mu^*) = h^{t(x, \mu^*)}(x, \mu^*) - \max_{t \neq t(x, \mu^*)} h^t(x, \mu^*) \leq \hat{h}_n^{t(x, \mu^*)}(x, \mu^*) - h^t(x, \mu^*) \quad \forall t \in \mathcal{T} \setminus \{t(x, \mu^*)\}, \forall x \in \mathcal{X}.$$

Equation (EC.19) implies that for a fixed  $\omega \in \hat{\Omega}$  and  $x \in \mathcal{X}$ , if  $\delta(x, \mu^*) > 0$ , then for  $\epsilon > 0$  small enough (and therefore  $n$  large enough),  $2\epsilon - \delta(x, \mu^*) < 0$ . This in turn implies  $\hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) - \hat{h}_n^{t(x, \mu^*)}(x, \hat{\mu}_n^*(\omega), \omega) < 0$  for all  $\forall t \in \mathcal{T} \setminus \{t(x, \mu^*)\}$ . Since  $t(x, \mu^*)$  is also the unique maximizer of  $\arg \max_{t \in \mathcal{T}} h^t(x, \mu)$  when  $\delta(x, \mu^*) > 0$ , this implies that the allocation of a sample-based dual-price queuing policy for an individual with covariates  $x$  match that of a dual-price queuing policy for  $n$  large enough for any  $\omega \in \hat{\Omega}$ , where  $\mathbb{P}(\hat{\Omega}) = 1$ . By Assumption 2 and Lemma EC.1, we have  $\delta(X, \mu^*) > 0$  almost surely. We will use this observation in Step 2 to prove that the expected average outcome of a sample-based dual-price policy almost surely will be as high as that of a true dual-price policy as  $n \rightarrow \infty$ .

**Step 2.** By Step 1, for any  $\omega \in \hat{\Omega}$  and sequence  $(\hat{\mu}_n^*(\omega))_{n \in \mathbb{N}}$  of sample-based solutions, where  $\hat{\mu}_n^*(\omega) \in \hat{\mathcal{S}}_n^*(\omega)$  for all  $n \in \mathbb{N}$ , there exists a sequence  $(\mu_n^*(\omega))_{n \in \mathbb{N}}$  of solutions, where  $\mu_n^*(\omega) \in \mathcal{S}^*$  for all  $n \in \mathbb{N}$ , with the following property. For every  $\epsilon > 0$ , there exists  $N(\omega) \in \mathbb{N}$  such that

$$\begin{aligned} \hat{h}_n^t(x, \hat{\mu}_n^*(\omega), \omega) - \hat{h}_n^{t(x, \mu_n^*(\omega))}(x, \hat{\mu}_n^*(\omega), \omega) < 2\epsilon - \delta(x, \mu_n^*(\omega)) \\ \forall t \in \mathcal{T} \setminus \{t(x, \mu_n^*(\omega))\}, \forall x \in \mathcal{X}, \forall n \geq N(\omega). \end{aligned} \quad (\text{EC.20})$$

Fix a  $\omega \in \hat{\Omega}$  and sequence couple  $(\hat{\mu}_n^*(\omega))_{n \in \mathbb{N}}$  and  $(\mu_n^*(\omega))_{n \in \mathbb{N}}$  such that (EC.20) holds. From now on, we slightly abuse the notation and drop  $\omega$  notation. In the following, intuitively, we fix a path of historical samples, which implies that our estimates  $\hat{\mu}_n^*$ ,  $\hat{m}_n^t$ ,  $\hat{h}_n^t$  and  $\hat{z}_n^D$  are no longer random. We still perceive the covariates and outcomes of the individuals arriving during the implementation as random and will denote by  $\mathbb{E}_x$  the expectation taken with respect to the distribution of these individuals' covariates. Recalling that  $z^D = z^*$  by Theorem EC.1, we have

$$\begin{aligned}
\limsup_{n \rightarrow \infty} z^* - \hat{z}_n^D &= \limsup_{n \rightarrow \infty} z^D - \hat{z}_n^D \\
&\leq \limsup_{n \rightarrow \infty} \mathbb{E}_x \left[ C \mathbb{1} \left\{ \min \arg \max_{t \in \mathcal{T}} h^t(X, \mu_n^*) \neq \min \arg \max_{t \in \mathcal{T}} \hat{h}_n^t(X, \hat{\mu}_n^*) \right\} \right] \\
&\leq \mathbb{E}_x \left[ \limsup_{n \rightarrow \infty} C \mathbb{1} \left\{ \min \arg \max_{t \in \mathcal{T}} h^t(X, \mu_n^*) \neq \min \arg \max_{t \in \mathcal{T}} \hat{h}_n^t(X, \hat{\mu}_n^*) \right\} \right] \\
&\leq \mathbb{E}_x \left[ \limsup_{n \rightarrow \infty} C \mathbb{1} \left\{ \hat{h}_n^t(X, \hat{\mu}_n^*) - \hat{h}_n^{t(X, \mu_n^*)}(X, \hat{\mu}_n^*) \geq 0 \text{ for some } t \in \mathcal{T} \setminus \{t(X, \mu_n^*)\} \right\} \right] \\
&\leq \mathbb{E}_x \left[ \limsup_{n \rightarrow \infty} C \mathbb{1} \left\{ \hat{h}_n^t(X, \hat{\mu}_n^*) - \hat{h}_n^{t(X, \mu_n^*)}(X, \hat{\mu}_n^*) + \delta(X, \mu_n^*) > 0 \text{ for some } t \in \mathcal{T} \setminus \{t(X, \mu_n^*)\} \right\} \right], \tag{EC.21}
\end{aligned}$$

where the first inequality holds because a gap between  $z^D$  and  $z_n^D$  occurs only when the dual-price queuing policy and its sample-based variant assign different treatments to an individual with covariates  $X$  and because  $m^t$  is uniformly bounded by  $C$  by Assumption 2 (i), which implies that the gap is also bounded by the same value. The second inequality follows from the reverse Fatou's lemma, which applies because the indicator function is bounded above by 1. The third inequality holds because for every  $x \in \mathcal{X}$ ,  $\min \arg \max_{t \in \mathcal{T}} h^t(x, \mu_n^*) \neq \min \arg \max_{t \in \mathcal{T}} \hat{h}_n^t(x, \hat{\mu}_n^*)$  implies that there exists a  $t \in \mathcal{T} \setminus \{t(x, \mu_n^*)\}$  such that  $\hat{h}_n^t(x, \hat{\mu}_n^*) - \hat{h}_n^{t(x, \mu_n^*)}(x, \hat{\mu}_n^*) \geq 0$ , where  $t(x, \mu_n^*)$  is defined as before, i.e.,  $t(x, \mu_n^*) = \min \arg \max_{t \in \mathcal{T}} h^t(x, \mu_n^*)$ . To see this, for a fixed  $x \in \mathcal{X}$ , if  $t(x, \mu_n^*) \notin \arg \max_{t \in \mathcal{T}} \hat{h}_n^t(x, \hat{\mu}_n^*)$ , then there exists  $t \in \mathcal{T} \setminus \{t(x, \mu_n^*)\}$  such that  $\hat{h}_n^t(x, \hat{\mu}_n^*) - \hat{h}_n^{t(x, \mu_n^*)}(x, \hat{\mu}_n^*) > 0$ . On the other hand, if  $t(x, \mu_n^*) \in \arg \max_{t \in \mathcal{T}} \hat{h}_n^t(x, \hat{\mu}_n^*)$  but is not the minimum index treatment, then there exists  $t \in \mathcal{T} \setminus \{t(x, \mu_n^*)\}$  such that  $\hat{h}_n^t(x, \hat{\mu}_n^*) - \hat{h}_n^{t(x, \mu_n^*)}(x, \hat{\mu}_n^*) = 0$ . Finally, the last inequality holds because  $\hat{h}_n^t(X, \hat{\mu}_n^*) - \hat{h}_n^{t(X, \mu_n^*)}(X, \hat{\mu}_n^*) \geq 0$  implies that  $\hat{h}_n^t(X, \hat{\mu}_n^*) - \hat{h}_n^{t(X, \mu_n^*)}(X, \hat{\mu}_n^*) + \delta(X, \mu_n^*) > 0$  almost surely as  $\delta(X, \mu_n^*) > 0$  almost surely by Assumption 2 and Lemma EC.1.

We will next show that the last expression in (EC.21) equals zero. By (EC.20), we have

$$\limsup_{n \rightarrow \infty} \hat{h}_n^t(x, \hat{\mu}_n^*) - \hat{h}_n^t(x, \hat{\mu}_n^*) + \delta(x, \mu_n^*) \leq 0 \quad \forall t \in \mathcal{T} \setminus \{t(x, \mu_n^*)\}, x \in \mathcal{X}. \tag{EC.22}$$

Equation (EC.22) implies that for any  $x \in \mathcal{X}$ , we also have

$$\liminf_{n \rightarrow \infty} \mathbb{1} \left\{ \hat{h}_n^t(x, \hat{\mu}_n^*) - \hat{h}_n^t(x, \hat{\mu}_n^*) + \delta(x, \mu_n^*) \leq 0 \quad \forall t \in \mathcal{T} \setminus \{t(x, \mu_n^*)\} \right\} = 1.$$

We then have

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \mathbb{1} \left\{ \hat{h}_n^t(x, \hat{\mu}_n^*) - \hat{h}_n^{t(x, \mu_n^*)}(x, \hat{\mu}_n^*) + \delta(x, \mu_n^*) > 0 \text{ for some } t \in \mathcal{T} \setminus \{t(x, \mu_n^*)\} \right\} \\
&= 1 - \liminf_{n \rightarrow \infty} \mathbb{1} \left\{ \hat{h}_n^t(x, \hat{\mu}_n^*) - \hat{h}_n^{t(x, \mu_n^*)}(x, \hat{\mu}_n^*) + \delta(x, \mu_n^*) \leq 0 \quad \forall t \in \mathcal{T} \setminus \{t(x, \mu_n^*)\} \right\} \\
&= 0.
\end{aligned}$$

As the above equality holds true for all  $x \in \mathcal{X}$ , the last expression in (EC.21) therefore amounts to zero. By (EC.21), we thus have  $\limsup_{n \rightarrow \infty} z^* - \hat{z}_n^D \leq 0$  for the fixed path of historical samples. The claim now follows from the observation that our choice of sequence couple  $(\hat{\mu}_n^*)_{n \in \mathbb{N}}$  and  $(\mu_n^*)_{n \in \mathbb{N}}$  was arbitrary and that  $\limsup_{n \rightarrow \infty} z^* - \hat{z}_n^D \leq 0$  holds for every sequence couple  $(\hat{\mu}_n^*)_{n \in \mathbb{N}} = (\hat{\mu}_n^*)_{n \in \mathbb{N}}(\omega)$  and  $(\mu_n^*)_{n \in \mathbb{N}} = (\mu_n^*)_{n \in \mathbb{N}}(\omega)$ , where  $\omega \in \hat{\Omega}$  and  $\mathbb{P}(\hat{\Omega}) = 1$ . In other words, we have  $\limsup_{n \rightarrow \infty} z^* - \hat{z}_n^D \leq 0$  almost surely.  $\square$

### EC.3. Derivations of Fairness-Constrained Policies for Section 4

#### EC.3.1. Allocation-Parity-Constrained Policies

In this section, we show the derivation of policy  $\hat{\pi}_{\text{alloc},n}^*$  of Section 4.1. The derivation for the minority prioritization case is similar and omitted. Recall that we now explicitly consider the protected feature  $G$  as an input into the policy  $\pi$ , i.e.,  $X = (X^{-G}, G)$ . The sample approximation of the allocation-parity-constrained policy design problem is given by

$$\begin{aligned}
\hat{z}_{\text{alloc},n}^* &= \max_{\pi \in \Pi} \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) \hat{m}_n^t(x_i) \\
\text{s.t.} \quad & \frac{1}{n} \sum_{i=1}^n \pi^t(x_i) \leq b^t \quad \forall t \in \mathcal{T} \\
& \frac{1}{n_g} \sum_{i=1}^n \mathbb{1}[g_i = g] \pi^t(x_i) - \frac{1}{n_{g'}} \sum_{i=1}^n \mathbb{1}[g_i = g'] \pi^t(x_i) \leq \delta \quad \forall g, g' \in \mathcal{G}, g \neq g', t \in \mathcal{T},
\end{aligned} \tag{EC.23}$$

where  $n_g$  denotes the number of individuals with protected feature  $g$  in the data. Let  $\mu \in \mathbb{R}_+^{m+1}$  collect the dual variables of the capacity constraints and denote by  $\lambda^t(g, g') \in \mathbb{R}_+$  the dual variable of the fairness constraint for pair  $(g, g')$  and for treatment  $t$ . The Lagrangian dual of problem (EC.23) is given by

$$\hat{v}_{\text{alloc},n}^* = \min_{\substack{\mu \in \mathbb{R}_+^{m+1} \\ \lambda \in \mathcal{L}_\infty(\mathcal{G} \times \mathcal{G}, \mathbb{R}_+^{m+1})}} \max_{\pi \in \Pi} L_{\text{alloc},n}(\pi, \mu, \lambda), \tag{EC.24}$$

where

$$\begin{aligned}
L_{\text{alloc},n}(\pi, \mu, \lambda) &= \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) \hat{m}_n^t(x_i) + \sum_{t \in \mathcal{T}} \mu^t \left( b^t - \frac{1}{n} \sum_{i=1}^n \pi^t(x_i) \right) + \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda^t(g, g') \delta \\
&\quad - \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda^t(g, g') \left( \frac{1}{n_g} \sum_{i=1}^n \mathbb{1}[g_i = g] \pi^t(x_i) - \frac{1}{n_{g'}} \sum_{i=1}^n \mathbb{1}[g_i = g'] \pi^t(x_i) \right).
\end{aligned}$$

We can rewrite the last term as

$$\begin{aligned} & \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda^t(g, g') \left( \frac{1}{n_g} \sum_{i=1}^n \mathbb{1}[g_i = g] \pi^t(x_i) - \frac{1}{n_{g'}} \sum_{i=1}^n \mathbb{1}[g_i = g'] \pi^t(x_i) \right) \\ &= \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G}} \left( \frac{1}{n_g} \sum_{i=1}^n \mathbb{1}[g_i = g] \pi^t(x_i) \right) \left( \sum_{g' \in \mathcal{G}, g \neq g'} (\lambda^t(g, g') - \lambda^t(g', g)) \right). \end{aligned}$$

For ease of notation, we let  $\gamma^t(g) = \sum_{g' \in \mathcal{G}, g \neq g'} (\lambda^t(g, g') - \lambda^t(g', g))$ . We now have

$$\begin{aligned} L_{\text{alloc},n}(\pi, \mu, \lambda) &= \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) \hat{m}_n^t(x_i) + \sum_{t \in \mathcal{T}} \mu^t \left( b^t - \frac{1}{n} \sum_{i=1}^n \pi^t(x_i) \right) \\ &\quad + \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda^t(g, g') \delta - \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G}} \left( \frac{1}{n_g} \sum_{i=1}^n \mathbb{1}[g_i = g] \pi^t(x_i) \right) \gamma^t(g) \\ &= \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) \left( \hat{m}_n^t(x_i) - \mu^t - \frac{n}{n_{g_i}} \gamma^t(g_i) \right) + \sum_{t \in \mathcal{T}} \mu^t b^t + \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda^t(g, g') \delta. \end{aligned}$$

For any fixed  $\mu$  and  $\lambda$ , the inner maximization problem of (EC.24) is solved by

$$\pi^t(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} \left( \hat{m}_n^{t'}(x) - \mu^{t'} - \frac{n}{n_g} \gamma^{t'}(g) \right) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T}, x \in \mathcal{X}. \quad (\text{EC.25})$$

We suppress the dependence of the above policy  $\pi$  on  $\mu, \lambda$  notationally in order to avoid clutter and tie-break using a lexicographic tie-breaker. By substituting the above policy into the dual problem (EC.24), we obtain the convex program

$$\hat{v}_{\text{alloc},n}^* = \min_{\substack{\mu \in \mathbb{R}_+^{m+1}; \\ \lambda \in \mathcal{L}_\infty(\mathcal{G} \times \mathcal{G}, \mathbb{R}_+^{m+1})}} \bar{L}_{\text{alloc},n}(\mu, \lambda),$$

where

$$\bar{L}_{\text{alloc},n}(\mu, \lambda) = \frac{1}{n} \sum_{i=1}^n \max_{t \in \mathcal{T}} \left( \hat{m}_n^t(x_i) - \mu^t - \frac{n}{n_{g_i}} \gamma^t(g_i) \right) + \sum_{t \in \mathcal{T}} \mu^t b^t + \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda^t(g, g') \delta.$$

We denote by  $\hat{\mu}_n^*, \hat{\lambda}_n^*$  an optimal solution to this finite convex program. We can now define our allocation-parity-constrained policy  $\hat{\pi}_{\text{alloc},n}^*$  by substituting  $\hat{\mu}_n^*, \hat{\lambda}_n^*$  into (EC.25).

### EC.3.2. Outcome-Parity-Constrained Policies

We now derive the policy  $\hat{\pi}_{\text{out},n}^*$  of Section 4.2. The sample approximation of the outcome-parity-constrained policy design problem is given by

$$\begin{aligned} \hat{z}_{\text{out},n}^* &= \max_{\pi \in \Pi} \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) \hat{m}_n^t(x_i) \\ \text{s.t.} \quad & \frac{1}{n} \sum_{i=1}^n \pi^t(x_i) \leq b^t \quad \forall t \in \mathcal{T} \\ & \frac{1}{n_g} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \mathbb{1}[g_i = g] \pi^t(x_i) \hat{m}_n^t(x_i) - \frac{1}{n_{g'}} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \mathbb{1}[g_i = g'] \pi^t(x_i) \hat{m}_n^t(x_i) \leq \delta \\ & \quad \forall g, g' \in \mathcal{G}, g \neq g'. \end{aligned} \quad (\text{EC.26})$$

Again let  $\mu \in \mathbb{R}_+^{m+1}$  be the dual variable for the capacity constraints, and denote by  $\lambda(g, g') \in \mathbb{R}_+$  the dual variable of the statistical parity in outcome constraint for the group pair  $g$  and  $g'$ . The Lagrangian relaxation of problem (EC.26) is given by

$$\hat{v}_{\text{out},n}^* = \min_{\substack{\mu \in \mathbb{R}_+^{m+1} \\ \lambda \in \mathcal{L}_\infty(\mathcal{G} \times \mathcal{G}, \mathbb{R}_+)}} \max_{\pi \in \Pi} L_{\text{out},n}(\pi, \mu, \lambda), \quad (\text{EC.27})$$

where

$$\begin{aligned} L_{\text{out},n}(\pi, \mu, \lambda) = & \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) \hat{m}_n^t(x_i) + \sum_{t \in \mathcal{T}} \mu^t \left( b^t - \frac{1}{n} \sum_{i=1}^n \pi^t(x_i) \right) + \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda(g, g') \delta \\ & - \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda(g, g') \left( \frac{1}{n_g} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \mathbb{1}[g_i = g] \pi^t(x_i) \hat{m}_n^t(x_i) - \frac{1}{n_{g'}} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \mathbb{1}[g_i = g'] \pi^t(x_i) \hat{m}_n^t(x_i) \right). \end{aligned}$$

We can rewrite the last term as

$$\begin{aligned} & \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda(g, g') \left( \frac{1}{n_g} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \mathbb{1}[g_i = g] \pi^t(x_i) \hat{m}_n^t(x_i) - \frac{1}{n_{g'}} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \mathbb{1}[g_i = g'] \pi^t(x_i) \hat{m}_n^t(x_i) \right) \\ & = \sum_{g \in \mathcal{G}} \left( \frac{1}{n_g} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \mathbb{1}[g_i = g] \pi^t(x_i) \hat{m}_n^t(x_i) \right) \left( \sum_{g' \in \mathcal{G}, g \neq g'} (\lambda(g, g') - \lambda(g', g)) \right). \end{aligned}$$

For ease of notation, we let  $\gamma(g) = \sum_{g' \in \mathcal{G}, g' \neq g} (\lambda(g, g') - \lambda(g', g))$ . We now have

$$L_{\text{out},n}(\pi, \mu, \lambda) = \frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \pi^t(x_i) \left( \hat{m}_n^t(x_i) \left( 1 - \frac{n}{n_{g_i}} \gamma(g_i) \right) - \mu^t \right) + \sum_{t \in \mathcal{T}} \mu^t b^t + \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda(g, g') \delta.$$

For any fixed  $\mu$  and  $\lambda$ , the inner maximization problem of (EC.27) is thus solved by

$$\pi^t(x) = \begin{cases} 1 & \text{if } t = \min \arg \max_{t' \in \mathcal{T}} \left( \hat{m}_n^{t'}(x) \left( 1 - \frac{n}{n_g} \gamma(g) \right) - \mu^{t'} \right) \\ 0 & \text{otherwise} \end{cases} \quad \forall t \in \mathcal{T}, x \in \mathcal{X}. \quad (\text{EC.28})$$

We suppress the dependence of policy  $\pi$  on  $\mu, \lambda$  notationally and tie-break using a lexicographic tie-breaker. By substituting the above policy into the dual problem (EC.27), we obtain the convex program

$$\hat{v}_{\text{out},n}^* = \min_{\substack{\mu \in \mathbb{R}_+^{m+1} \\ \lambda \in \mathcal{L}_\infty(\mathcal{G} \times \mathcal{G}, \mathbb{R}_+)}} \bar{L}_{\text{out},n}(\mu, \lambda),$$

where


$$\bar{L}_{\text{out},n}(\mu, \lambda) = \frac{1}{n} \sum_{i=1}^n \max_{t \in \mathcal{T}} \left( \hat{m}_n^t(x_i) \left( 1 - \frac{n}{n_{g_i}} \gamma(g_i) \right) - \mu^t \right) + \sum_{t \in \mathcal{T}} \mu^t b^t + \sum_{g \in \mathcal{G}} \sum_{g' \in \mathcal{G}, g \neq g'} \lambda(g, g') \delta.$$

We denote by  $\hat{\mu}_n^*, \hat{\lambda}_n^*$  an optimal solution to this convex program. We define the outcome-parity-constrained policy  $\hat{\pi}_{\text{out},n}^*$  by substituting  $\hat{\mu}_n^*, \hat{\lambda}_n^*$  into (EC.28).

## EC.4. Empirical Results Details

### EC.4.1. Linear Outcomes Simulated Data

In Figure EC.1, we show additional simulation results for values of  $\sigma$  from the set  $\{0.5, 0.8, 1.15\}$  for the noise term. The qualitative takeaways are the same as those in Section 5.1.3 and serve primarily to show the degradation of non-parametric methods of decision tree and KNN from added noise in the data.



**Figure EC.1** Synthetic data results for  $\sigma$  values  $\{0.5, 0.8, 1.15\}$ . All subfigures show the ratio of out-of-sample performance of the sample-based queuing policy to that of the perfect foresight policy in dependence of training set size. Each line corresponds to a policy using a different outcome estimator.

### EC.4.2. Nonlinear Outcomes Simulated Data

In this section, we investigate a variant of the setup from Section 5.1.1, where the outcomes are now nonlinear in the covariates. In particular, the potential outcomes  $Y^t$  and conditional mean treatment outcomes  $m^t$ , for each  $t \in \{0, 1, 2\}$ , can be written as:

$$Y^0 = m^0(X) + \epsilon_0, \quad m^0(X) = 0.25X_1^2 + 0.75X_2^2$$

$$Y^1 = m^1(X) + \epsilon_1, \quad m^1(X) = 0.75X_1^2 + 0.75X_2^2$$

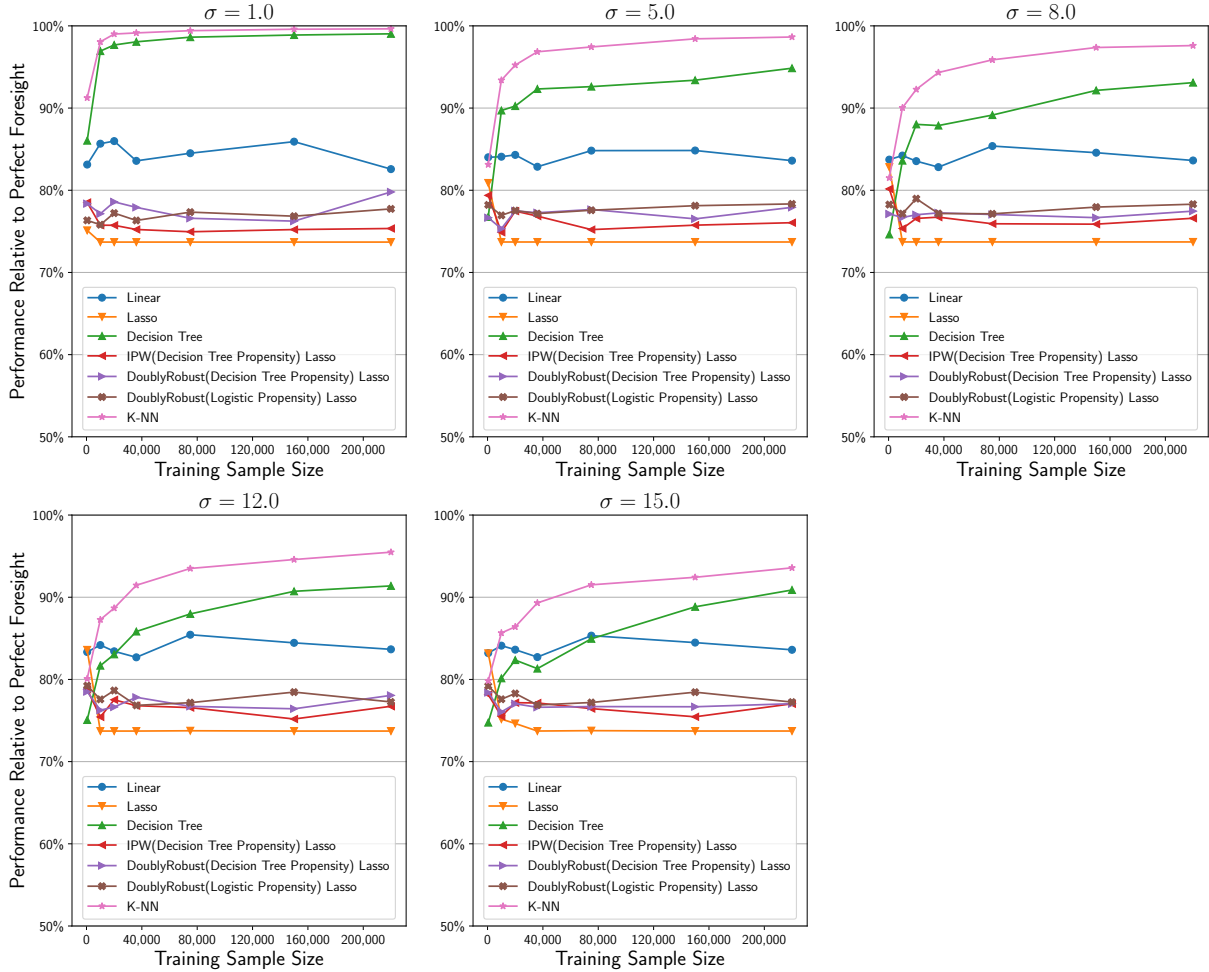
$$Y^2 = m^2(X) + \epsilon_2, \quad m^2(X) = 0.25X_1^2 + 1.25X_2^2$$

where  $\epsilon_t \sim \mathcal{N}(0, \sigma)$ ,  $t \in \mathcal{T}$ , are i.i.d. noise terms. We now vary  $\sigma$  over a larger range of values  $\{1, 5, 8, 12, 15\}$  since the quadratic outcome functions  $m^0, m^1, m^2$  now take on a larger range of values (using the range of  $\sigma$  values from the linear case of Section 5.1.1 does not yield noise terms with meaningful differences compared to the scale of outcome values).

Figure EC.2 shows the simulation results for this setting. Compared to the case of linear conditional mean treatment outcomes, as one would expect, no policy based on a linear outcome estimator reaches close to 100% of the perfect performance, regardless of sample size, since they



are all misspecified now. In the low noise regime ( $\sigma = 1$ ), the policy based on the non-parametric  $k$ -NN estimator comes close to 100% (as does using decision trees), but increasing levels of noise cause a deterioration in the performance of  $k$ -NN. We also see that simpler parametric models like linear and lasso, though misspecified, are relatively robust to increasing levels of noise. The causal adjusted versions of lasso also improve their performance as well here, though only to a small degree.



**Figure EC.2** Synthetic data results for nonlinear outcome function for  $\sigma$  values  $\{1, 5, 8, 12, 15\}$ . All subfigures show the ratio of out-of-sample performance of the sample-based queuing policy to that of the perfect foresight policy in dependence of training set size. Each line corresponds to a policy using a different outcome estimator.

### EC.4.3. HMIS and VI-SPDAT Data

**EC.4.3.1. Data Details and Preparation** In this subsection, we provide further details on the current system, raw data, and how we defined outcomes for our experimental results.

When an individual arrives to the system, they are assessed for vulnerability with the VI-SPDAT tool, which contains a series of questions such as prior length of and number of experiences of homelessness. Adverse answers, such as sleeping situations indicative of homelessness, are given a weight of 1 so that an individual’s vulnerability score ranges from 0 to 17, where a higher score indicates more vulnerability. The question and answers from the tool are used as features in our estimation methods and a sample of the tool is shown in Figure EC.3 (LAHSA 2017). To determine treatment assignments and outcomes after individuals receive their VI-SPDAT assessments, we can match them with their enrollment history from the raw HMIS data, which details their interactions with the system. Full details on the HMIS data and variables can be found in the HMIS Data Standards Data Dictionary (U.S. HUD 2021). For a given individual, we observe a sequence of enrollments, each characterized by a type of service and the corresponding date information.

<b>A. History of Housing and Homelessness</b>			
<b>4. Where do you sleep most frequently?</b>		<input type="checkbox"/> Shelters <input type="checkbox"/> Transitional Housing <input type="checkbox"/> Safe Haven <input type="checkbox"/> Outdoors* <input type="checkbox"/> Couch surfing	<input type="checkbox"/> Car <input type="checkbox"/> Client doesn't know* <input type="checkbox"/> Client refused* <input type="checkbox"/> Other* (please specify: _____ )
		<b>If the person answers anything other than "Shelters", "Transitional Housing", or "Safe Haven", then score 1.</b>	<b>Score:</b>
<b>5. How long has it been since you lived in permanent stable housing?</b>	<input type="checkbox"/> Less than a week <input type="checkbox"/> 1 week – 3 months <input type="checkbox"/> 3 – 6 months	<input type="checkbox"/> 6 months to 1 year <input type="checkbox"/> 1 – 2 years <input type="checkbox"/> 2 years or more	<input type="checkbox"/> Client doesn't know <input type="checkbox"/> Client refused
<b>6. In the last three years, how many times have you been homeless?</b>	<input type="checkbox"/> 0 times <input type="checkbox"/> 1 time <input type="checkbox"/> 2 times	<input type="checkbox"/> 3 times <input type="checkbox"/> 4 times <input type="checkbox"/> 5 or more times	<input type="checkbox"/> Client doesn't know <input type="checkbox"/> Client refused
		<b>If the person has experienced 1 or more consecutive years of homelessness, and/or 4+ episodes of homelessness, then score 1.</b>	<b>Score:</b>

**Figure EC.3 VI-SPDAT sample**

Finally note that there is not a one-to-one correspondence between individuals found in VI-SPDAT assessment data and the HMIS data. After merging the two datasets, and removing missing data issues, we obtain our final dataset that consists of 63,764 individuals assessed between 1/12/2015 and 12/31/2019, their assessment information, treatment received, and observed outcome. Recall that we use all individuals assessed between 1/12/2015 to 12/31/2017 as our training set to learn treatment outcomes and construct our policy, and then evaluate on individuals assessed between 1/1/2018 to 12/31/2019 to measure out-of-sample performance.

*Historical Treatment Assignment Definition.* Since the original HMIS data does not explicitly contain “Treatment”, we use the enrollment history to create these variables. In Table EC.1 below, we breakout the 11 types of enrollments found within HMIS data, where the enrollments ‘PH-PSH’ and ‘PH-RRH’ are considered permanent interventions or treatments, and the remaining enrollments are considered no-treatment, or ‘Services Only’, either due to their temporary nature, shelter type, or nature of service provided. We also further breakout ‘PH-PSH’ into ‘tenant’ and ‘site’ based on an additional column, where the two sub-types differ, broadly speaking, in terms of supportive services available on-site and process of actually receiving housing. This categorization of the enrollment types was chosen based on discussions with matchers working within the LA CES system.

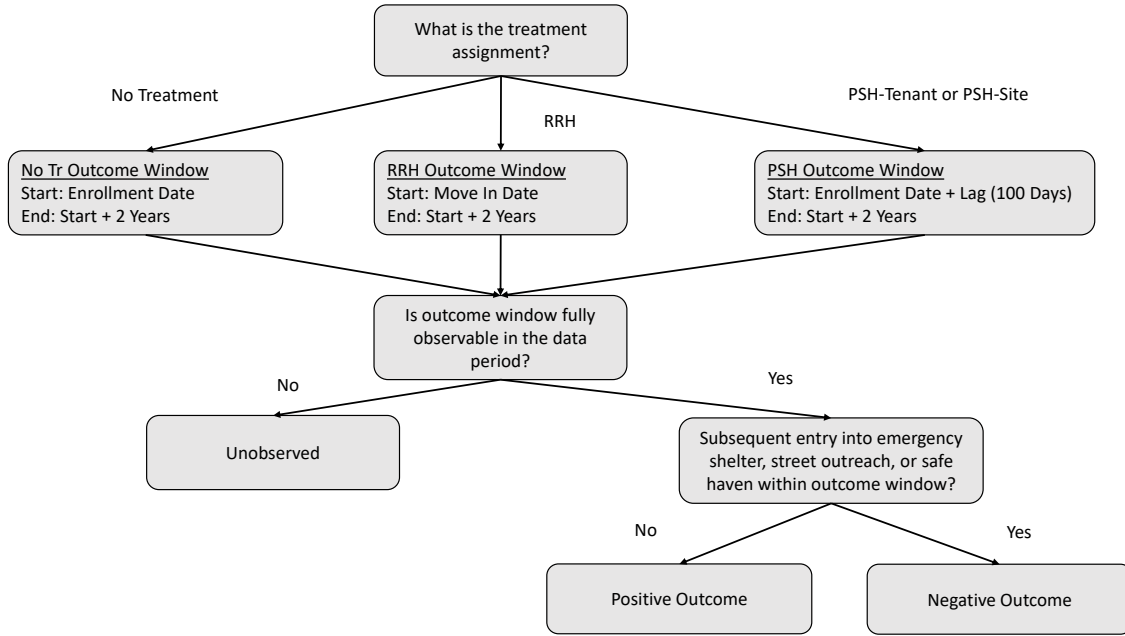
**Table EC.1** HMIS enrollment types grouped by enrollments considered as permanent interventions and no treatments. Percentage breakdowns come from our data used for estimation and are based on individuals who received a VI-SPDAT assessment.

Treatment Type	Enrollment Type	Proportion of All Enrollments
Permanent Intervention	PH - RRH	14.16%
Permanent Intervention	PH - PSH Site	4.63%
Permanent Intervention	PH - PSH Tenant	4.14%
No Treatment	Emergency Shelter	26.74%
No Treatment	Street Outreach	26.23%
No Treatment	Services	18.73%
No Treatment	Other	3.10%
No Treatment	Transitional Housing	1.31%
No Treatment	Day Shelter	0.55%
No Treatment	Homelessness Prevention	0.30%
No Treatment	PH - Other	0.09%
No Treatment	Safe Haven	0.05%

*Outcome Definition.* Recall from Section 5.2.1 that we focus on measuring returns to homelessness after an intervention as the outcome variable, in line with the goals listed by HUD. Since such an outcome is not explicitly tracked within the data, we construct a proxy variable as follows. To determine an individual’s outcome, we first define an two-year observation window depending on their treatment received. For those receiving ‘PH-RRH’, we also observe a ‘Move-In Date’ column signifying when an individual was recorded to have officially received the resource. We define the

‘PH-RRH’ window to start from the ‘Move-In Date’. Similarly for ‘PH-PSH’, we use the ‘Move-In Date’ to determine the starting point, but add a 100 day lag to the ‘Move-In Date’ because we concluded the recorded ‘Move-In Date’ within HMIS is likely not accurate. Through discussions with matchers within the LA CES system, we learned that the ‘Move-In Date’ within HMIS is an initial date inputted by case managers, but often does not get updated after an individual eventually moves in. Furthermore, we analyzed an auxiliary dataset, the Resource Management System (RMS) data, that supports this conclusion. For a subset of PSH units found within the HMIS data, the RMS data contains specific information such as their eligibility requirements and the updated move-in dates. For units found within both the HMIS and RMS data, we found discrepancies between the reported move-in dates, where the RMS ‘Move-In Date’ occurs often months after the corresponding HMIS ‘Move-In Date’, indicating that the HMIS ‘Move-In Date’ column for ‘PH-PSH’ have inaccuracies. This presents a problem where we often found individuals receiving PSH within the HMIS data to subsequently enroll in enrollments indicating an individual was experiencing homelessness (‘Street Outreach’, ‘Emergency Shelter’, and ‘Safe Haven’) shortly after their reported HMIS PSH ‘Move-In Date’. If the HMIS ‘Move-In Date’ was incorrect and those individuals had actually not yet moved-in yet to their PSH units, then we would incorrectly conclude they experienced an adverse outcome after receiving PSH.

To address the problems with the HMIS PSH ‘Move-In Date’ column, we add a 100 day lag to the ‘Move-In Date’ to determine the starting point for the ‘PH-PSH’ observation window. We determined this 100 day lag based on the distribution of positive differences between the RMS and HMIS ‘Move-In Date’. Finally for those receiving **no-treatment**, we consider their first enrollment date of any type to be the start of their observation window. Our choice of a two-year observation window for determining if an individual returned to homelessness after an intervention involved a trade-off where shorter windows resulted in insufficient time to observe post-intervention outcomes while longer windows resulted in fewer individuals for whom we can observe their full outcome window. If we observe any occurrence of the three proxies of homelessness enrollments (‘Street Outreach’, ‘Emergency Shelter’, or ‘Safe Haven’) within that observation window, then we consider an individual to have experienced a negative outcome. Otherwise we consider them to have a positive outcome. We made the above choices of treatment types, window length, start dates, and more considering the trade-offs in terms of number of samples, ‘bias’ of the proxies, and data quality. For individuals for whom we do not fully observe their outcome window, i.e., the outcome window ends after the last observation date in our data, we remove them from our analysis. Keeping those individuals would bias the data towards positive outcomes since individuals may experience negative outcomes after the observable time in our data. For a summary of our outcome column construction, see Figure EC.4.

**Figure EC.4 Construction of treatment outcomes.**

**EC.4.3.2. Generating Counterfactuals and Fitting Training Set Estimators** To obtain counterfactuals for evaluating performance within the held out test set, we took a semi-synthetic approach and used model-generated outcomes based on the entire dataset. We consider three causal estimation approaches for generating counterfactual outcomes, the direct (or T-Learner method), IPW, and DR methods. While the direct method simply fits an observed outcome estimator for each treatment group, the IPW and DR methods involve fitting an additional propensity model that estimates the probability of receiving each treatment conditional on observed covariates. The inverse propensity scores are then used as sample weights in a second stage counterfactual estimator in order to correct for selection bias. Since many causal estimation methods involve a two step procedure where the first step estimates nuisance components like outcome and propensity models and the second step combines these components into a final counterfactual model, recent works often refer to the final model as a “final” or “second” stage estimator (Künzel et al. 2019). Therefore, we also adopt this terminology going forward and refer to the final counterfactual estimator as a second stage estimator. We largely follow similar procedures as those described in Battocchi et al. (2019) and Jacob (2021). For counterfactual outcome model selection for each treatment, we want our model output probabilities to be as well-calibrated to the observed outcome distribution as possible. Since our outcome is binary, the probability of positive outcome for individual with features  $X$  under treatment  $t$  is exactly the function  $m^t(X)$ . Therefore, requiring calibrated probabilities under each treatment is the same as requiring  $\hat{m}^t(x) \sim m^t(x)$  for all  $x \in \mathcal{X}$  for our chosen treatment outcome generating model  $\hat{m}^t$ . Therefore, we generate counterfactuals

from models that empirically match the outcome distribution of each treatment as well as possible. To ensure generalization as well, we take a standard cross-validation approach to select the best model for each treatment group. Finally, we noticed individuals without any indication of a disability in the data were very unlikely to receive PSH and, generally speaking, PSH units required some form of disability to be eligible. Therefore modeling the PSH outcomes of non-disabled individuals would possibly violate the positivity assumption in Assumption 1 and practically we have insufficient samples. Therefore, we model disabled and non-disabled individuals separately and did not model non-disabled individual outcomes under PSH treatment.

Note that the procedure outlined in this section was applied to the entire dataset and does not relate to the time-based splitting of training (1/12/2015 to 12/31/2017 assessments) and testing sets (1/1/2018 to 12/31/2019 assessments) used to evaluate performance in our experiment. Here, we are simply trying to find models that empirically fit our data and outcome distributions well enough to serve as counterfactuals.

*Cross Validation.* Since we want to have well-calibrated probabilistic predictions, we use log-loss during the cross validation phase as the evaluation metric for tuning hyper-parameters of each model class for estimating observed outcomes, propensities, and any second stage estimator. We choose to minimize log-loss since log-likelihood is a strictly proper scoring rule, i.e., optimizing this metric will yield well-calibrated probability predictions (Gneiting and Raftery 2007). For estimating each treatment group’s outcomes, treatment assignment propensities, and second stage components, we use 5-fold cross validation and select the hyper-parameters with the best average log-loss across all folds for each model class we considered (logistic regression with and without regularization, random forests, gradient boosted trees, and XGBoost trees).

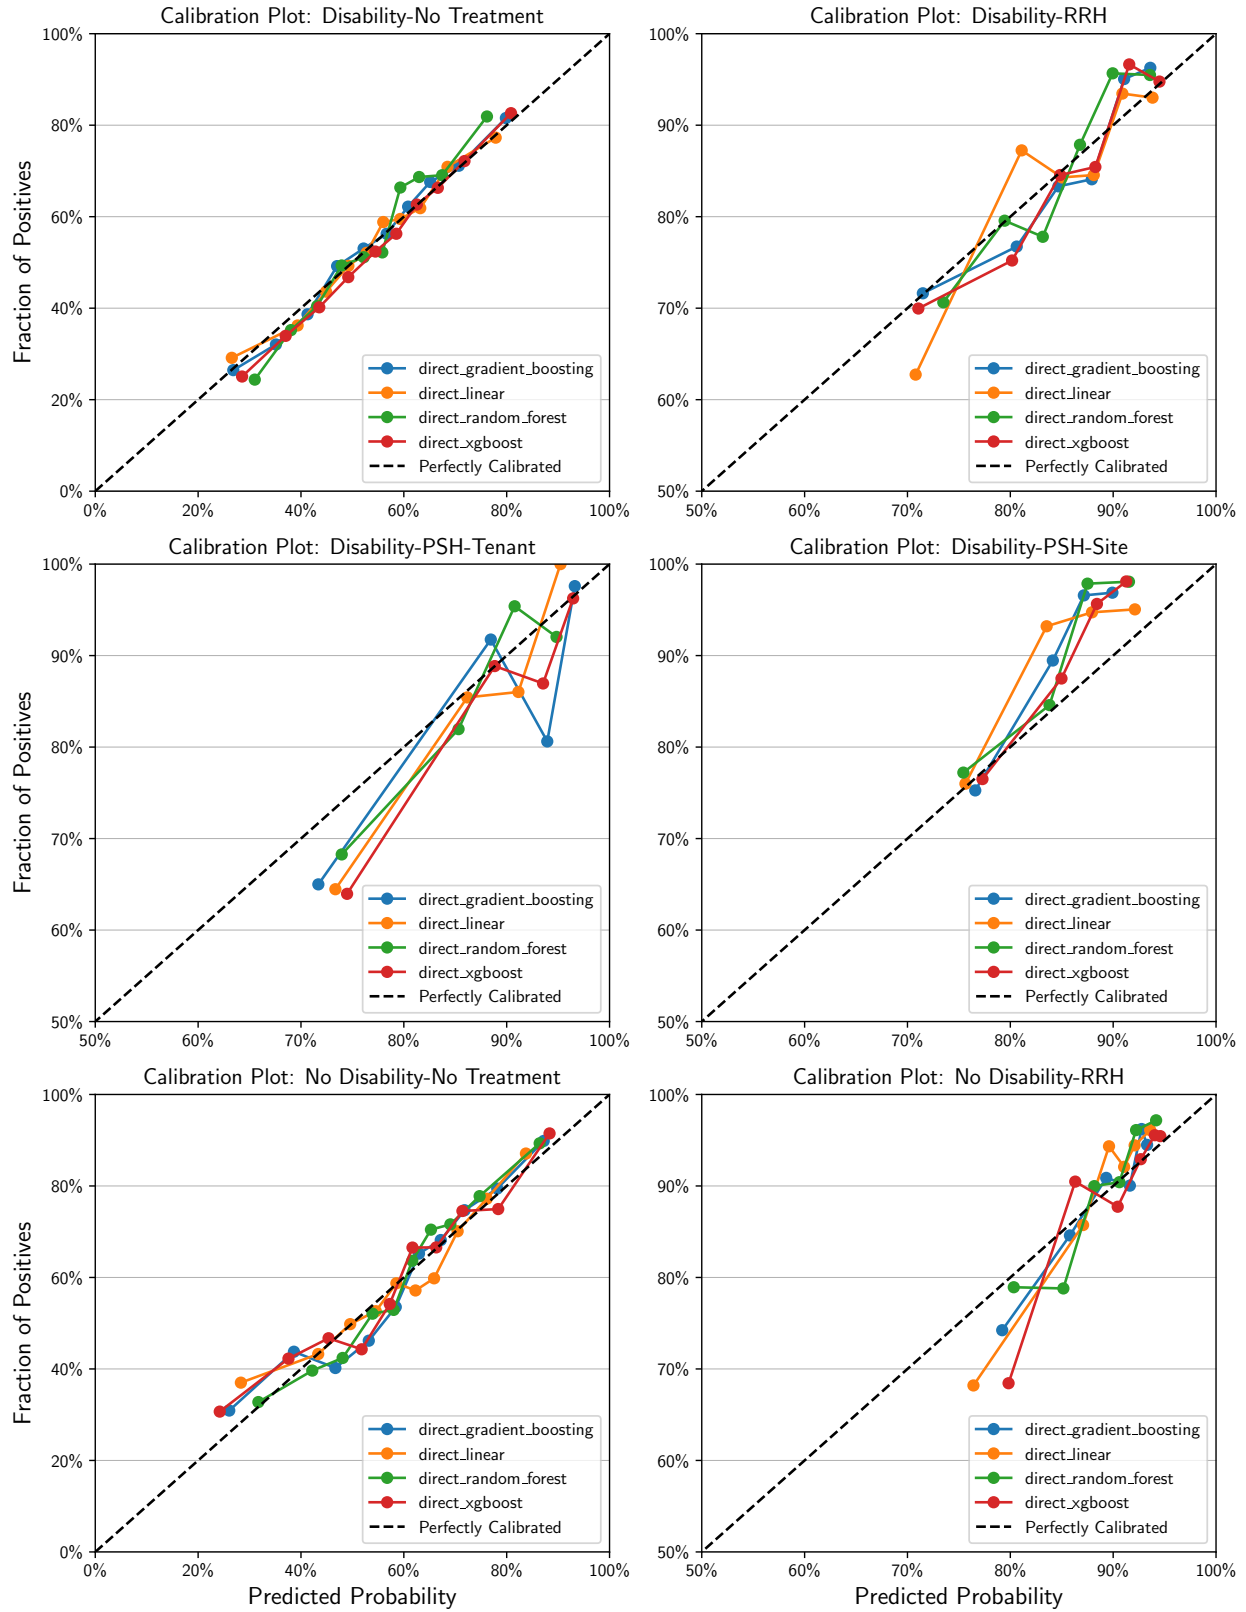
*Validation Set of Calibration.* For each treatment group and hyper-parameter tuned counterfactual outcome model, we generate a calibration plot to visually pick the model that is best calibrated. In Figures EC.5, EC.6, EC.7, we plot the calibration curves from a held-out validation set for the direct, IPW, and DR methods, respectively. For example, in each of the subplots of Figure EC.5, the dotted diagonal line represents probability outputs of a perfectly calibrated model while the colored solid lines with markers are the calibration curves of each model class. In particular, we calculate weighted calibration curves. To generate each calibration curve, we split the probability predictions for a given model into equal sized buckets and find the average predicted probability (x-axis) and plot it against the weighted average proportion of positive outcomes (y-axis), where the weights are the inverse propensity weights. In essence, because we can only plot and observe outcomes for those who received a particular treatment, we need to correct for observational data selection bias. One way to correct for selection bias is to calculate each bucket’s average outcome

using the IPW estimator (Hernan and Robins 2023). Also, because the treatment groups **RRH**, **PSH-Tenant**, and **PSH-Site** had much smaller numbers of samples relative to the **no-treatment** group, we used a smaller number of buckets to evaluate calibration to avoid each bucket being too small in size and resulting in noise. Generally speaking, there was not a big difference between the calibration curves of the direct and IPW methods, while the DR curves on the smaller treatment groups look more miscalibrated. The DR miscalibration may be due to the weighted residuals used in the doubly robust second stage estimation, as the residuals, which are functions of a binary outcome subtracting a predicted probability, may be better suited for debiasing continuous outcomes. Given that overall the best IPW calibration curve within each disability and treatment group appeared the most calibrated, we defaulted to using the IPW models for generating counterfactual outcomes.

After selecting the appropriate estimator for each treatment group of disability and non-disability groups, we refitted the chosen models on the training and validation sets and checked for out-of-sample generalization on the held-out test set. In Figure EC.8. we plot the final test-set calibration curves for each treatment group for individuals with and without disabilities. In general, the final chosen models for each treatment group appear well-calibrated visually with some slight deviations from the diagonal line. Finally, while the above procedure describes how we chose the counterfactual generating models, we still need to generate predictions of treatment outcomes for the training set only (all individuals assessed 1/12/2015 to 12/31/2017) to learn a sample-based dual-price queuing policy. We repeat the above procedure but only on the training data to generate  $\hat{m}^t$ .

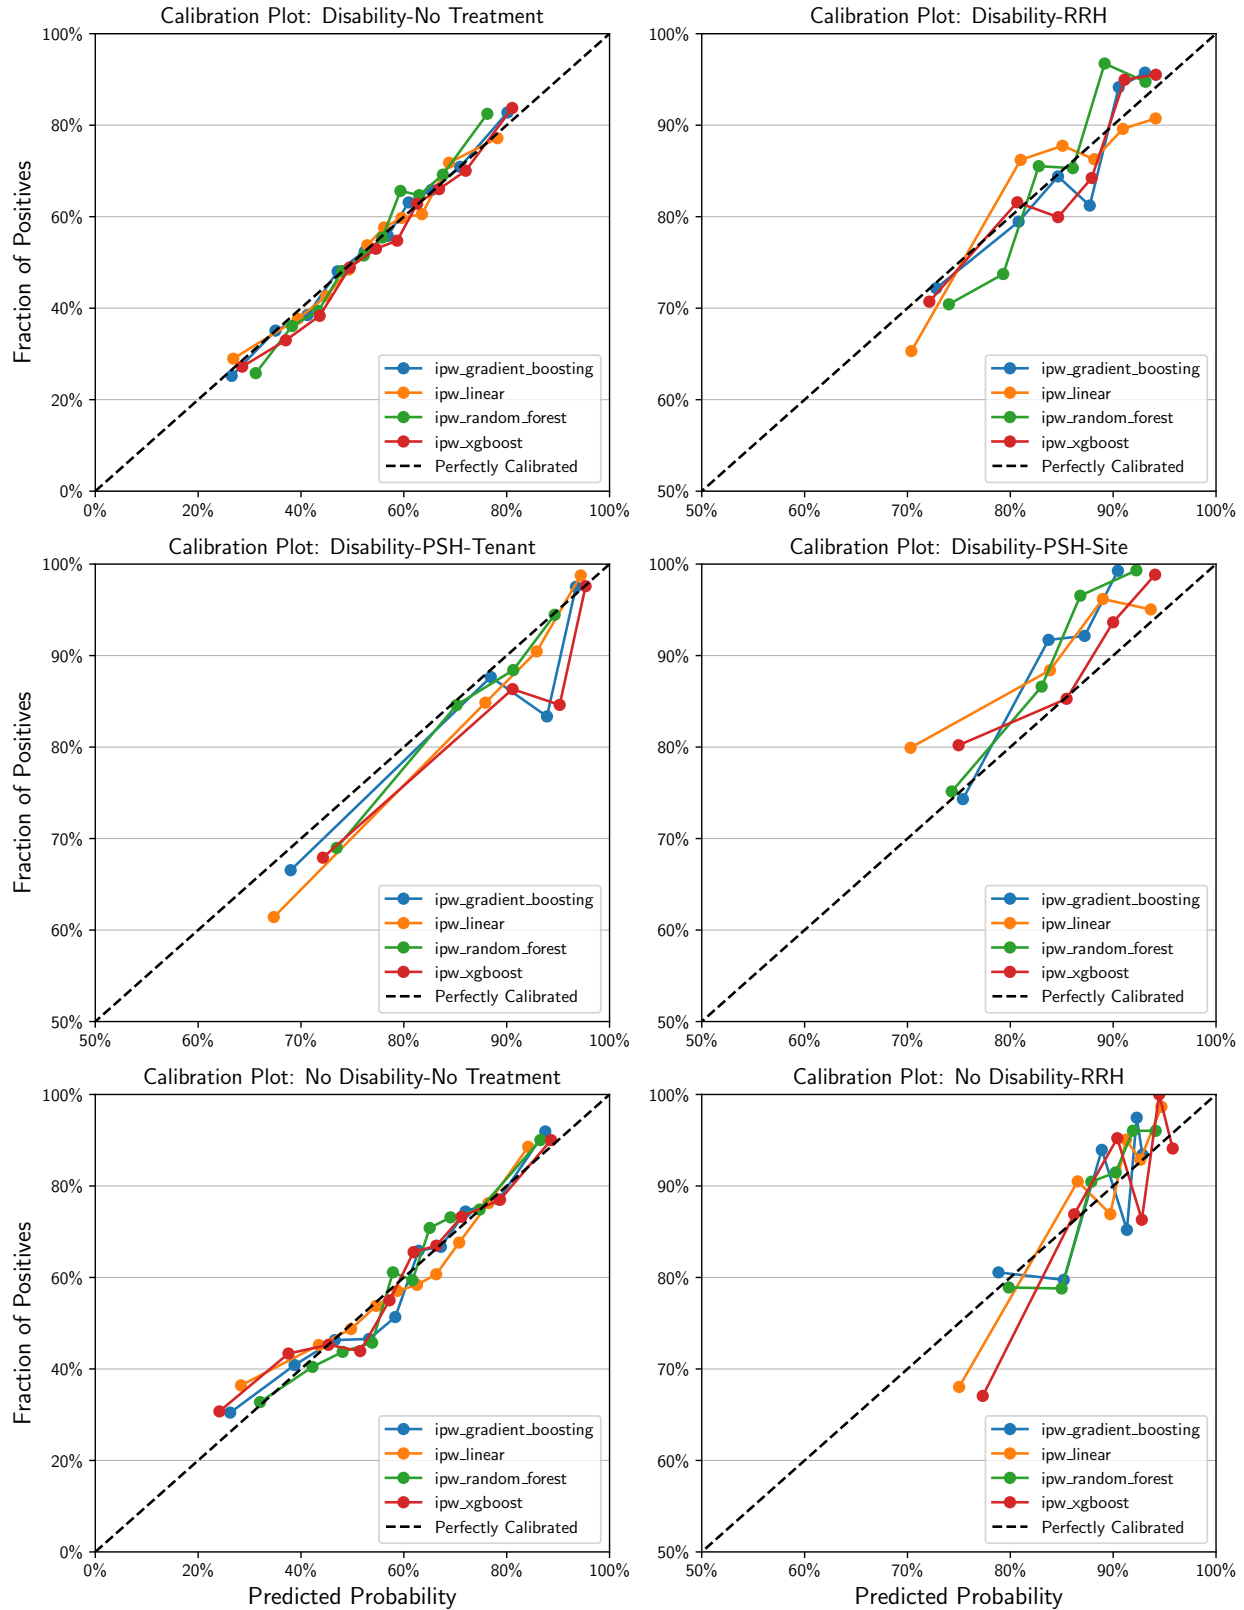
**EC.4.3.3. Fairness of Estimates** An additional concern with our counterfactual models for evaluation was potential bias by race. Since we want to use our generated counterfactuals as reflections of the true generating process, any calibration bias would mean our evaluation results would be potentially biased by race. In Figure EC.9, we plot and compare calibration curves for each racial group across treatments on the test-set for the final chosen models. For the most part, the calibration curves across racial groups are similar for each treatment, with the exception of the ‘Other’ racial group. This is likely because the ‘Other’ group is only a small portion of the sample, representing only  $\sim 7\%$  of the sample. Furthermore, the scarce **PSH-Tenant** and **PSH-Site** treatment groups also have worse calibration than **RRH**, again likely due to insufficient samples from this small treatment group.

**EC.4.3.4. Full Experimental Test Results** In this subsection, we present further results on two additional fairness-constrained policies **Allocation SP** and **Outcome SP** policies, which seek to achieve statistical parity in allocation (constraint (4)) and statistical parity in outcomes (constraint (7)) by race, respectively. The out-of-sample performance in terms of expected outcomes are summarized in Table EC.2 and allocation and outcome fairness properties are presented in

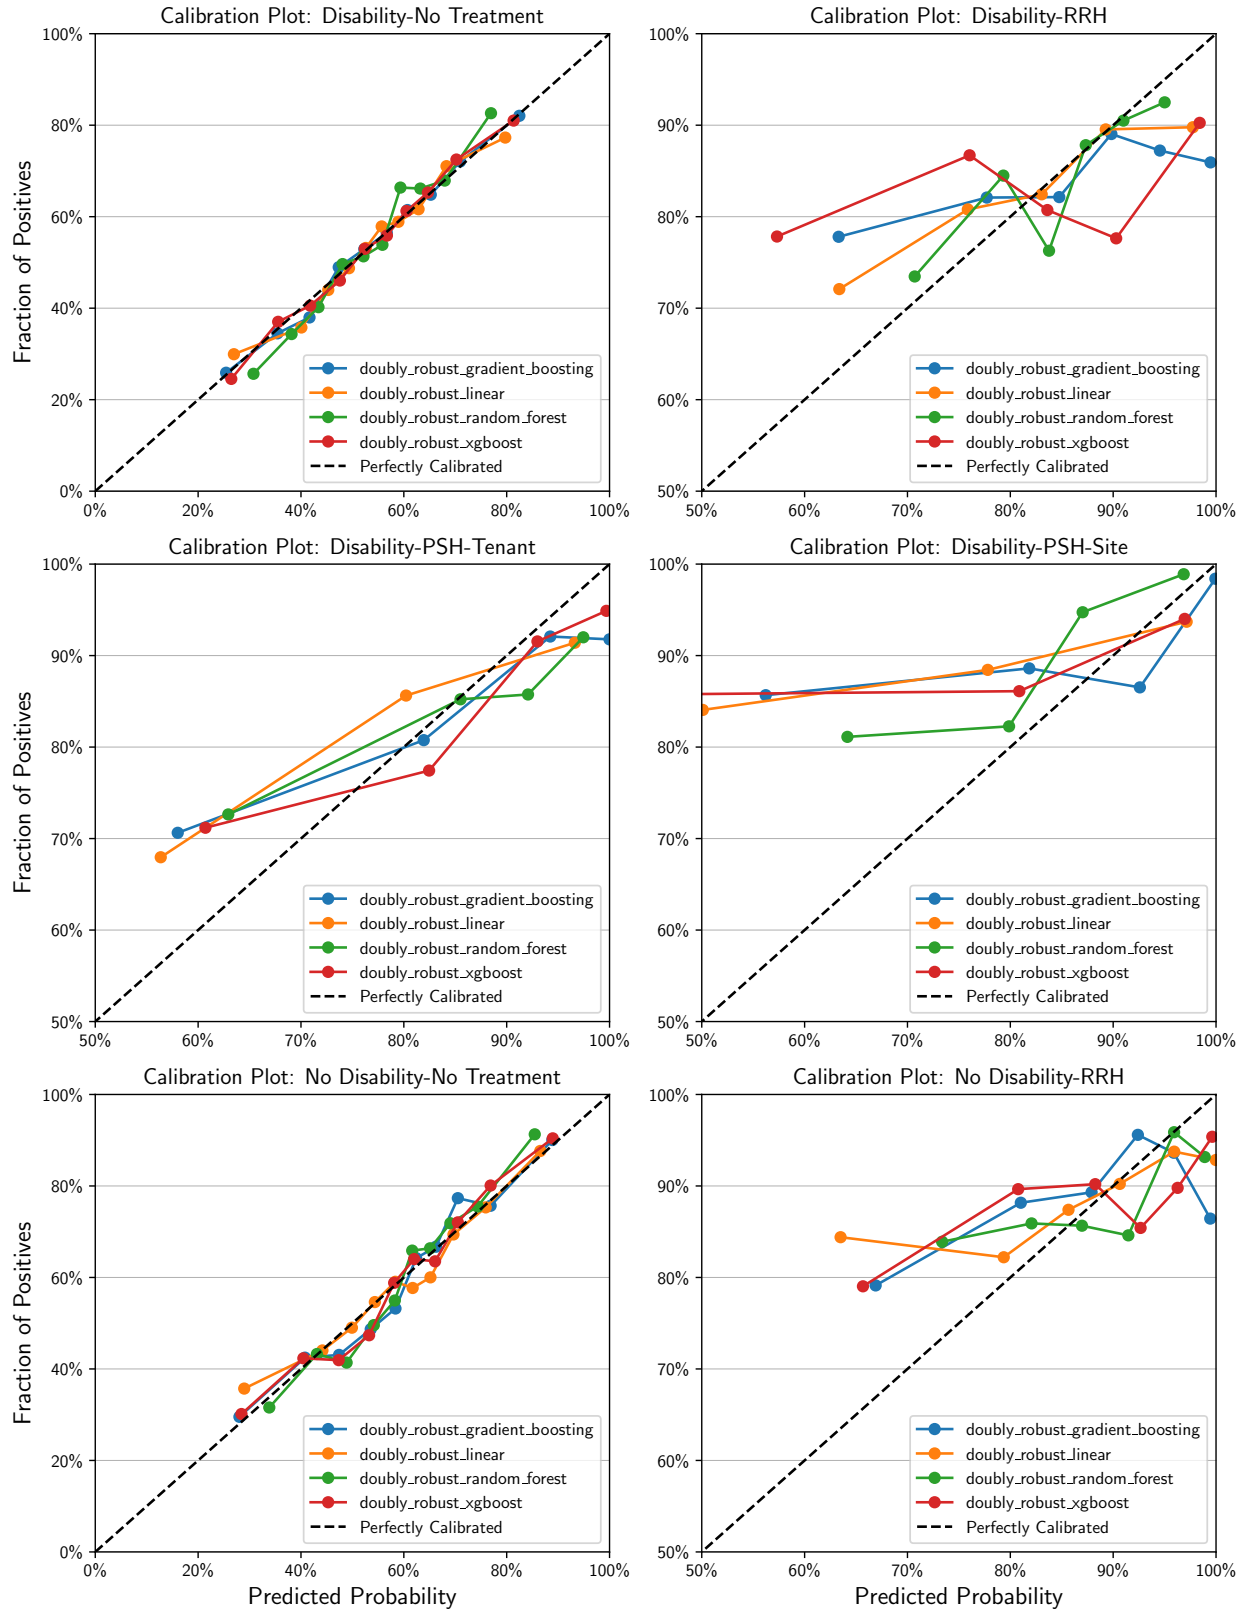


**Figure EC.5** Calibration plots on validation set for each direct method outcome model class and different combinations of disability and treatment group.

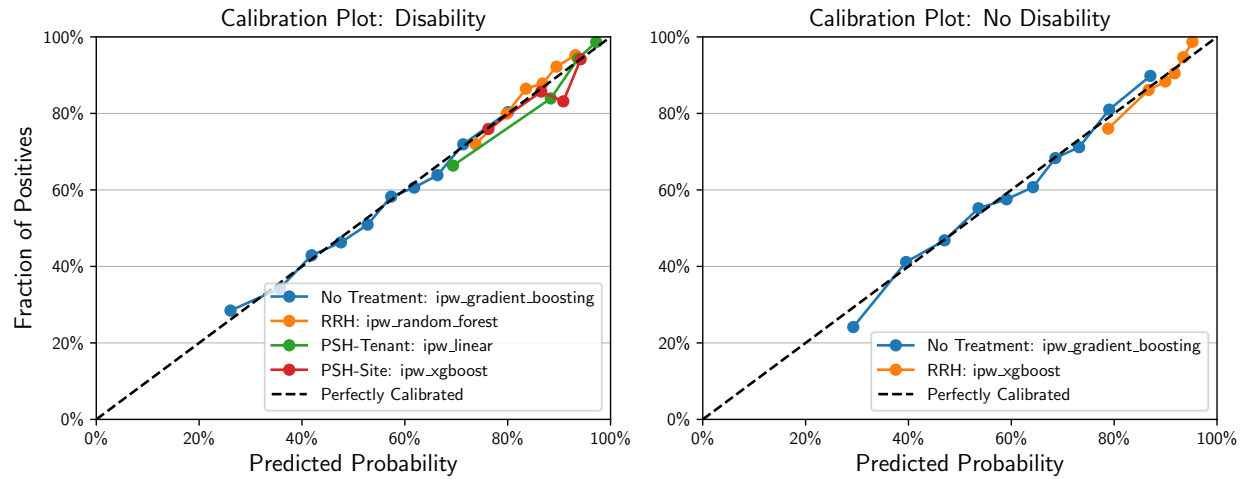




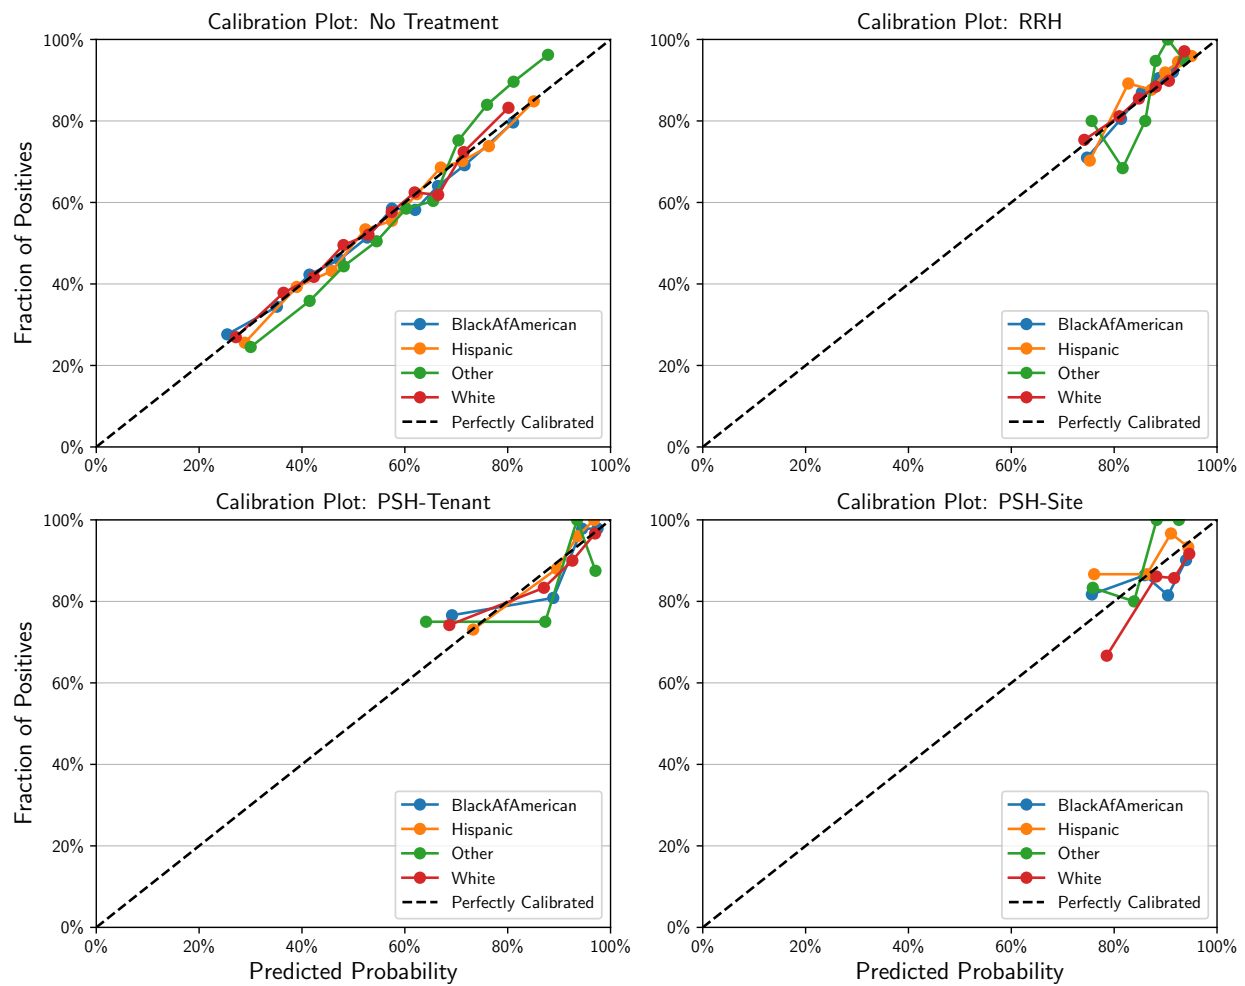
**Figure EC.6** Calibration plots on validation set for each IPW method outcome model class and different combinations of disability and treatment group.



**Figure EC.7** Calibration plots on validation set for each Doubly Robust method outcome model class and different combinations of disability and treatment group.



**Figure EC.8** Out-of-sample calibration plots of final counterfactual generating models chosen for each treatment group and for individuals with and without disabilities.



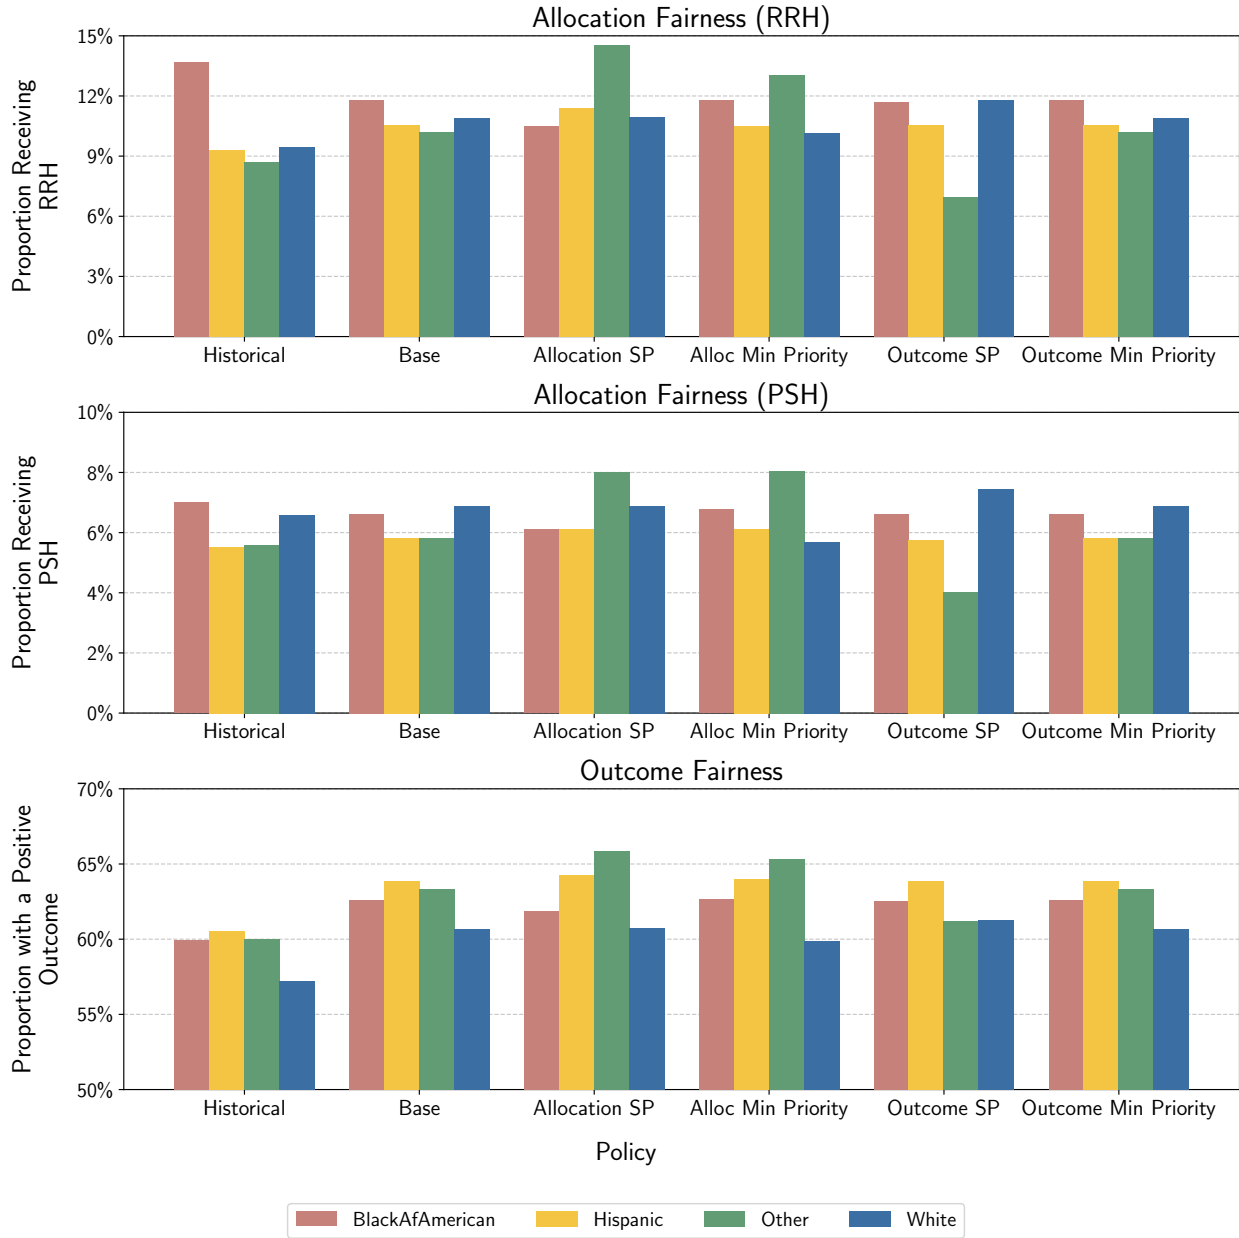
**Figure EC.9** Out-of-sample calibration plots for different racial groups and treatments from final counterfactual generating models chosen.

**Table EC.2** Out-of-sample average expected outcome under each policy, and the relative percentage difference compared to the Historical allocation.

Policy	Proportion of Positive Outcomes	Percentage Change vs Historical
No Treatment	54.25%	−8.76%
Historical	59.46%	0.00%
Base	62.53%	5.16%
Allocation SP	62.50%	5.11%
Alloc Min Priority	62.51%	5.13%
Outcome SP	62.52%	5.15%
Outcome Min Priority	62.53%	5.16%
Perfect Foresight	63.45%	6.71%

Figure EC.10. We use a  $\delta$  of 0.01 for the fairness constraints in the sample-based problem to learn the Allocation SP and Outcome SP policies.

For RRH, we see that Allocation SP has the highest minimum percentage receiving RRH across all racial groups compared to all other policies, but surprisingly does not result in the smallest differences in percentages between all racial groups out-of-sample. It also allocates fewer resources to Black individuals compared to Historical or any other policy, which may be undesirable when policymakers want to prioritize minority groups. In this instance, policymakers need to choose between statistical parity in allocation or prioritizing minority groups for allocation of RRH but our proposed policy can incorporate either notions into consideration or possibly some combination of the two. Similarly, for PSH, Allocation SP also has the highest minimum allocation rate across all racial groups, but again does not have the smallest differences between all racial groups and allocates to Black individuals at the lowest rate relative to all other policies. Also, of note in the allocation fairness results is that the Outcome SP policy allocates RRH and PSH to the Other group at far lower rates compared to other racial groups. This is because in our predictions, individuals of the Other group had predicted outcomes under no-treatment that skewed more positive relative to other racial groups. Therefore, the Outcome SP policy allocated less to the Other group in an attempt to achieve outcome parity across all groups because their predicted outcomes under no-treatment were already higher compared to the other racial groups. In terms of statistical parity in outcomes, we see that Outcome SP, in addition to Base and Outcome Min Priority, resulted in the smallest differences in proportion of positive outcomes between all racial groups out-of-sample.



**Figure EC.10** Out-of-sample fairness results showing proportion of each racial group receiving RRH, PSH (Allocation Fairness) and the proportion of positive outcomes for each racial group (Outcome Fairness) under different policies.

Finally, we see in Table EC.2 that our fairness-constrained policies **Allocation SP** and **Outcome SP** suffer almost no performance gap in terms of expected outcomes compared to **Base** while potentially improving the fairness properties of the **Base** policy. Again, these further results suggest almost no ‘price of fairness’.


#### EC.4.4. Qualitative Behavior of Dual-Price Queuing Policy

**EC.4.4.1. Allocations with Respect to Vulnerability** One potential concern for our policy is that by optimizing for overall expected outcomes, it may overlook vulnerable individuals who are not responsive to treatments. Under our queue assignment mechanism for the **Base** policy (Definition 1), each person is assigned to the queue  $t$  that maximizes the difference  $\hat{m}_n^t(x) - \hat{\mu}_n^{*,t}$ . Since the dual price for **No Treatment** is 0 (it has no budget constraint), a resource queue  $t$  is eligible for assignment only when  $\hat{m}_n^t(x) - \hat{\mu}_n^{*,t} - \hat{m}_n^0(x)$  is at least larger than 0. In this case, optimizing for expected outcomes will risk deprioritizing vulnerable individuals with low treatment effectiveness. However, we find this potential concern does not occur in our experimental results due to a strong correlation (93%) between an individual’s vulnerability and the treatment effectiveness of the best resource for that individual. This means that highly vulnerable individuals also tend to benefit most from treatment and are prioritized by our policy.

In Figure EC.11, we plot the probability of being matched to any treatment queue under the **Base** policy as a function of individual vulnerability, where vulnerability is defined as the probability of an adverse outcome under **No Treatment**. We observe a clear pattern where individuals with higher vulnerability are increasingly likely to receive a resource, because of their higher treatment effects. In particular, more vulnerable individuals are also more likely to receive the more supportive PSH resources. Therefore, in settings where individuals who are very vulnerable also tend to have large benefits from receiving treatments optimizing for expected outcomes does not necessarily come at the expense of vulnerable individuals.

**EC.4.4.2. Comparing Dual-Price Queueing Assignments to Historical Assignments Based on Characteristics** In practice, policymakers will want to understand why our policy allocates resources differently from the existing **Historical** policy. Since our policy consistently makes resource assignments as a function of individual characteristics, we can examine individual covariates to understand cases where our **Base** policy makes different assignments than the **Historical** policy. We therefore examined every individual in our test data and isolated the cases where the two policies disagreed: either **Base** offered *any* treatment while **Historical** offered none, or vice-versa. (Observations on which both policies agreed were discarded.)

Due to the high-dimensionality of our data (more than 100 covariates), we used a simple feature-importance analysis approach to understand the disagreements. For individuals with policy disagreements, we assigned a label 1 if **Base** assigned a treatment but **Historical** did not, and 0 in the other case. Then we fit a logistic regression with one-hot encoded categorical covariates and normalized continuous covariates so that all covariates were in the same scale. Finally, we rank covariates by the absolute value of their coefficients as a heuristic for feature importance.



**Figure EC.11** Proportion of individuals matched to any treatment queue under a dual-price queueing policy, as a function of individual vulnerability. Vulnerability is defined as the probability of an adverse outcome under no treatment, where the x-axis corresponds to buckets with a width of 10%. For example, 0% corresponds to the interval  $[0\%, 10\%)$ , 10% to  $[10\%, 20\%)$ , and so on, where each bucket is left-inclusive and right-exclusive. The y-axis represents the average matching probability for individuals within each vulnerability bucket.

The results show that the covariates with the highest absolute value coefficients for predicting whether **Base** vs **Historical** assigned an treatment were covariates that did not come directly from the VI-SPDAT survey but from auxiliary HMIS data sources. These included characteristics like an individual’s current employment type/status, whether they were looking for employment, or the specific category of living situation an individual was in. As an example, the **Historical** policy tended to assign treatments to individuals with some form of work whereas our policy was more likely to assign treatment to someone without known form of work. Since case managers primarily relied on the VI-SPDAT responses to make allocation decisions historically, our findings suggest that augmenting VI-SPDAT responses with additional HMIS data can improve allocation decisions.

## EC.5. Wait Time and Queue Length Statistics under the Dual-Price Queueing Policy

Section EC.5.1 discusses the theoretical properties of wait times and queue lengths under our proposed policy. Section EC.5.2 presents additional results on wait times observed in the numerical experiment described in Section 5.2.4. Section EC.5.3 investigates the causal effect of the wait times

observed under our counterfactual policy on outcomes. Section EC.5.4 provides further empirical evidence from the VI-SPDAT that an individual's circumstances are unlikely to materially worsen under our policy's wait times.

### EC.5.1. Discussion of Arrival and Queue Statistics

In this section, we argue that the expected number of individuals assigned to the queue for any resource under the dual-price queuing policy matches the expected number of resource arrivals of that type, ensuring balance in the system over time. In constructing our policy optimization framework, we have not made any assumptions regarding the arrival processes of individuals and resources in the system. For the purposes of this discussion, we now assume that the arrival process of individuals is such that the expected number of arrivals in any interval of length  $\ell$  is  $\lambda\ell$ , where  $\lambda > 0$ . Examples of such a process include cases where individuals arrive either with deterministic inter-arrival times of  $1/\lambda$  or with i.i.d. exponential inter-arrival times with mean  $1/\lambda$ , i.e., a Poisson arrival process. Similarly, we assume that the arrival process of resources of type  $t \in \mathcal{T} \setminus \{0\}$  is such that the expected number of arrivals in any interval of length  $\ell$  is  $\mu^t\ell$ , where  $\mu^t > 0$ . Note that the system is assumed to be initially empty. To reflect the scarcity of resources, we assume that  $\mu^t < \lambda$  for all  $t \in \mathcal{T} \setminus \{0\}$ .

In our policy optimization framework, the parameter  $b^t$ , which represents the asymptotic capacity of treatment  $t \in \mathcal{T} \setminus \{0\}$  per individual, can be related to the arrival processes through the equation  $b^t = \mu^t/\lambda$ , which denotes the expected number of arrivals of resources of type  $t$  per unit time divided by the expected number of individual arrivals. The policy  $\pi^*$ , defined in (EC.3), satisfies the capacity constraint  $\mathbb{E}[\pi^{*,t}(X)] \leq b^t$  for all  $t \in \mathcal{T}$ . Furthermore, under Assumption 2, if resources of type  $t \in \mathcal{T} \setminus \{0\}$  are very scarce and contribute to improving outcomes (which align with the housing application), these capacity constraints are binding for all  $t \in \mathcal{T} \setminus \{0\}$ .

The dual-price queueing policy introduced in Definition EC.1 assigns each individual with covariates  $x$  to the queue for treatment  $t$ , for  $t \in \mathcal{T}$ , if  $\pi^{*,t}(x) = 1$ . Given that the capacity constraints are binding for all  $t \in \mathcal{T} \setminus \{0\}$ , this implies that the expected number of individuals assigned to the queue for any treatment  $t \in \mathcal{T} \setminus \{0\}$  at any time matches the expected number of arrivals of resources of type  $t$ , assuming that the covariates  $X_k$  are i.i.d. across individuals. To see this, consider any time point  $\ell$  and any treatment type  $t \in \mathcal{T} \setminus \{0\}$ . Under the dual-price queueing policy, the expected number of individuals assigned to the queue for resource  $t$  at time  $\ell$  is given by  $b^t\lambda\ell$ . This follows from the fact that the prior probability of assigning an arriving individual to the queue for treatment  $t$  is  $\mathbb{E}[\pi^{*,t}(X)] = b^t$ , and the expected number of individuals arriving in the system until time  $\ell$  is  $\lambda\ell$ . Given the relation  $b^t = \mu^t/\lambda$ , this implies that the expected number of individuals assigned to the queue for resource  $t$  by time  $\ell$  matches the expected number of arriving resources



of type  $t$  up to the same time, i.e.,  $\mu^t \ell$ . In this sense, the dual-price queueing policy ensures a well-balanced system. If this balance is not maintained (i.e., if the prior probability of assigning a treatment  $t \in \mathcal{T} \setminus \{0\}$  is either low or high compared to that under  $\pi^*$ ), then, on expectation, the difference between the number of individuals waitlisted for treatment  $t$  and the number of resources of type  $t$  arriving in the system will continue to grow either toward infinity or negative infinity, indicating either longer waiting times for individuals or an increased number of unutilized resources—both of which are undesirable. So far, in the discussion, we have focused on the optimal dual-price queueing policy constructed under the assumption of knowing the distribution of the covariates and potential outcomes. Recall that, in reality, we only have access to data, meaning we can only compute and rely on the sample-based dual-price queueing policy, as defined in Definition 1. Since we have shown this policy to be asymptotically optimal, it is reasonable to expect that the sample-based dual-price queueing policy will achieve a similar balance as the number of in-sample data points grows.

Studying the statistics of queue lengths for any treatment  $t \in \mathcal{T} \setminus \{0\}$ , or the waiting times of individuals and resources in the system, is significantly more complex and depends on further specifications about the arrival processes as well as the distribution of covariates. To illustrate this complexity, consider the simplest case where inter-arrival times are deterministic; that is, at any time  $\ell$ ,  $\lambda \ell$  (assume to be an integer) individuals arrive in the system, and  $\mu^t \ell$  (assume to be an integer) resources of type  $t \in \mathcal{T} \setminus \{0\}$  become available. As the dual-price queueing policy assigns each arriving individual with covariates  $x$  to the queue for treatment  $t \in \mathcal{T} \setminus \{0\}$  with probability  $\pi^{*,t}(x)$ , the expected length of queue  $t$  at time  $\ell$  is given by

$$\mathbb{E} \left[ \left( \sum_{k=1}^{\lambda \ell} \pi^{*,t}(X_k) - \mu^t \ell \right)^+ \right], \quad (\text{EC.29})$$

where  $(a)^+ = \max\{a, 0\}$ . The term  $\sum_{k=1}^{\lambda \ell} \pi^{*,t}(X_k)$  represents the length of queue  $t$  after  $\lambda \ell$  arrivals of individuals with covariates  $X_1, \dots, X_{\lambda \ell}$ , while  $\mu^t \ell$  is the number of resources of type  $t$  that become available by time  $\ell$ . The expected queue length in equation (EC.29) depends not only on the expected allocation of an individual to queue  $t$  prior to knowing their covariates but also on the overall distribution of the covariates. Moreover, if we consider a stochastic arrival process, this quantity also depends on the specifics of that process (e.g., the distribution of inter-arrival times) and not merely on the expected number of arrivals, making the analysis significantly more complex.

Additionally, in reality, we only have access to data and therefore lack knowledge of the true distribution of the covariates as well as the specifics of the arrival process. When assessing queue lengths, waiting times, and the number of idle resources under our proposed sample-based queueing

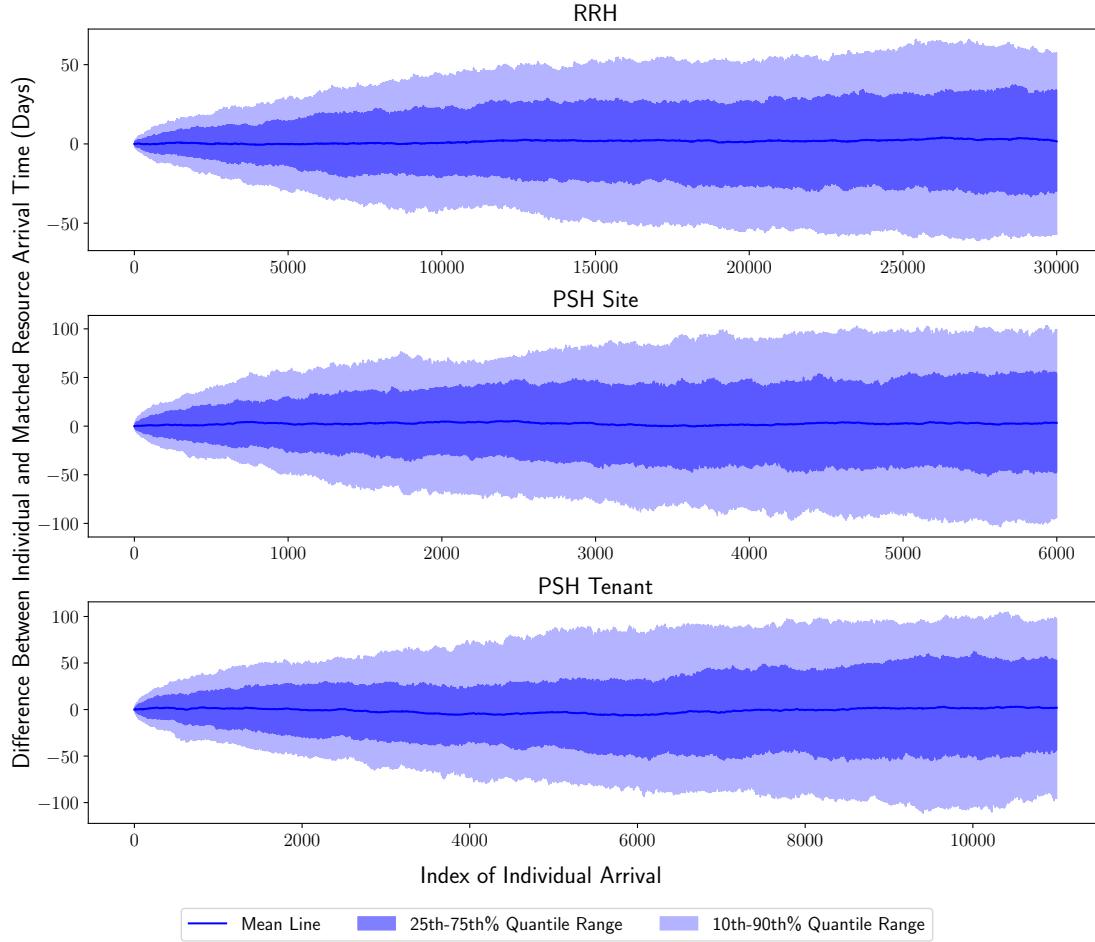
policy, we will therefore rely on data to estimate. In Section EC.5.2, we will indeed verify numerically that the allocation system under our sample-based queuing policy is well-balanced, in the sense that on average there is not a build up of individuals waiting in queues nor idle resources.

### EC.5.2. Full Wait Time Results for Section 5.2.4

In this subsection, we present further numerical results on wait times (Section EC.5.2.1) and queue lengths (Section EC.5.2.2) under our policy, based on the experiment setup described in Section 5.2.4.

**EC.5.2.1. Wait Times** Previously, we summarized our wait time results in Section 5.2.4 and Figure 3, which showed that our policy results in relatively low wait times on average, across 250 trials. One aspect that Figure 3 does not account for is idle housing “waiting” to be matched, which is also an undesirable situation based on our discussions with policy makers. Figure EC.12 presents the same statistics for the wait times of each  $i^{\text{th}}$  arriving individual, adjusted by subtracting the idle time of their matched resource. This “adjusted wait time” can be computed as the difference between the individual’s arrival time and the arrival time of their matched resource, and provides a holistic view of the allocation system in terms of both individuals and resources. The “adjusted wait time” is positive when an individual is waiting for a resource and is negative when a resource is waiting to be matched to an individual, and we clarify that individuals and resources cannot both be waiting. Note that the quantities in Figure 3 are obtained by truncating negative adjusted wait times, which represent idle resource times, to zero. We see from Figure EC.12 that across each treatment queue, the mean adjusted wait times stabilize around 0, indicating that the system is well-balanced on average, which aligns with the discussion in Section EC.5.1. For applications in which a holistic view of the allocation system is important, we may desire such well-balanced behavior, where on average there is neither waiting individuals nor idle resources. We note that the 10<sup>th</sup> quantiles (where adjusted wait times are negative and resources are idle) are  $-57$ ,  $-94$ , and  $-96$  days for RRH, PSH Site, or PSH Tenant, respectively. The 90<sup>th</sup> quantiles (where adjusted wait times are positive and individuals are waiting) are 57, 99, and 96 days for RRH, PSH Site, or PSH Tenant, respectively.

**EC.5.2.2. Queue Lengths** To complement the results on wait times, we additionally investigate the queue lengths under our policy based on our experiment setup. In Figure EC.13, we take a holistic system view by considering both individual and resource build-ups in treatment queues. We compute the queue size, displayed in this figure, as the difference between the number of individuals assigned and waiting for a particular housing type and the number of the respective resources waiting to be matched. Therefore, at each time point, a positive queue size means there are individuals waiting for resource arrivals while a negative queue size means there are resources

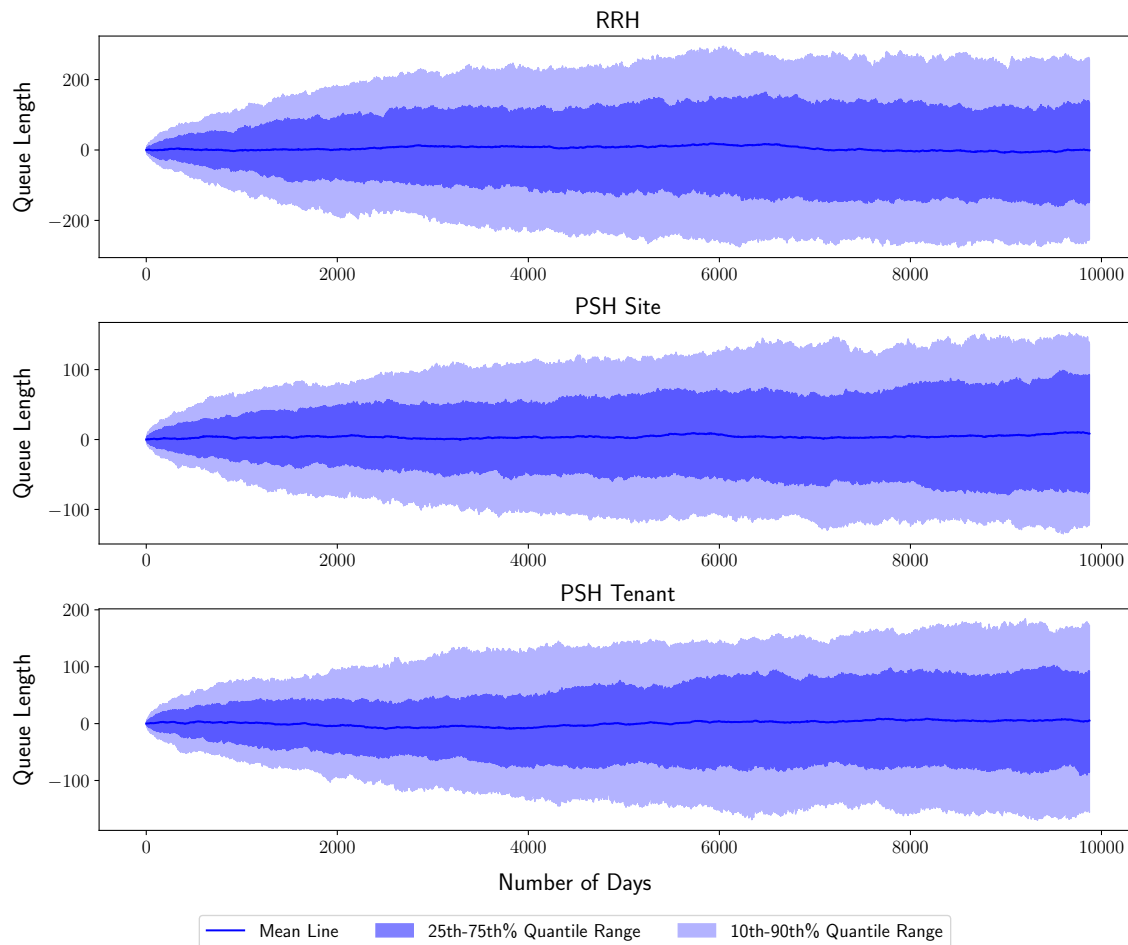


**Figure EC.12** Adjusted wait times for the different types of housing resources across 250 sample paths, where, for each individual, we plot the difference between an individual's arrival time and the arrival time of their matched resource. A positive value means an individual arrives before their matched resource while a negative value means a resource arrived before its matched individual.

waiting to be matched to an individual. Similar to the wait time results in Figure EC.12, we see that our policy, on average across simulation trials, results in a well-balanced system where neither individuals nor resources continuously build up within queues over time, but can result in longer queue lengths. The 10<sup>th</sup> quantiles (where queue sizes are negative and there are resources waiting) are -213, -118, and -156 across the subplots, respectively, and the 90<sup>th</sup> quantiles (where queue sizes are positive and individuals are waiting) are 256, 137, and 173 across the subplots, respectively.

### EC.5.3. Impact of Wait Time on Outcomes

One important consideration for whether the wait times observed under our policy are reasonable is the impact of these wait time levels on individual outcomes. Therefore, based on our historical data, we compute the negative impact of increasing wait times on individual outcomes for each



**Figure EC.13** Simulated treatment queue sizes across 250 simulated sample paths, where positive queue size means individuals are waiting in queue while negative queue size means resources are waiting to be matched.


group of individuals based on disability characteristic and resources received. To study potential heterogeneity in how wait times may impact different individuals, we compute wait time conditional average treatment effects (CATE)—treatment referring to wait time in this case—based on each individual’s covariates using the Orthogonal/Double Machine Learning (DoubleML) framework from the causal inference literature (Chernozhukov et al. 2018, Nie and Wager 2021). We chose this method because it allows for the incorporation of arbitrary machine learning techniques for estimation and naturally handles the continuous nature of wait times, in contrast to the binary or discrete treatment settings (recall that treatment refers to wait time in this analysis) commonly studied by other methods. For estimating the two nuisance functions required in DoubleML of outcomes and treatment assignment (in this case wait time) conditional on covariates and the final stage regression in DoubleML, we follow standard cross-validation and hyperparameter search procedures with the base estimators being linear/logistic regression (with and without regularization), random forests, gradient boosted trees, and XGBoost trees.

In Figure EC.14, we show the mean, 25<sup>th</sup> to 75<sup>th</sup>, and 10<sup>th</sup> to 90<sup>th</sup> quantile ranges of the decrease in probability of positive outcomes as wait times increase, where each subplot corresponds to a different combination of disability and resource received. In this case, a lower quantile (such as 10<sup>th</sup>%) corresponds to larger decreases in positive outcome probability as wait times increases, while higher quantile corresponds to individuals where wait time has less effect on their positive outcome probability. For the sample path average RRH wait time of 18 days, we see that the mean decrease in positive outcome probability is  $-0.53$  p.p. and  $-0.29$  p.p. for disabled and non-disabled individuals, respectively, while the 10<sup>th</sup> percentile decreases by  $-0.86$  p.p. and  $-0.61$  p.p. for disabled and non-disabled individuals, respectively. For PSH Tenant and PSH Site, the sample path average wait time of 30 days corresponds to mean decreases of  $-0.64$  and  $-0.36$  p.p. respectively, while the 10<sup>th</sup> percentile individual's decrease in positive outcome probability are  $-0.99$  and  $-0.87$  p.p., respectively. Even at 90<sup>th</sup> percentile wait times of 55, 100, and 100 days of wait time for RRH, PSH Tenant, and PSH Site, the mean decrease in positive outcome probabilities are  $-1.62$ ,  $-2.12$ ,  $-1.21$ , and  $-0.90$  p.p. for Disabled RRH, PSH Tenant, PSH Site, and Non-Disabled RRH individuals respectively. In addition, at 90<sup>th</sup> percentile wait times, the 10<sup>th</sup> percentile individual's decrease in positive outcome probability are  $-2.62$ ,  $-3.31$ ,  $-2.89$ , and  $-1.85$  p.p. across the subplots, respectively. As expected, increasing wait times has a negative effect on positive outcome probabilities across individuals, but the impact is relatively minimal at the wait times seen in Figure 3.

#### EC.5.4. Impact of Time on Vulnerability: Evidence from VI-SPDAT Reassessments

Another way to gauge the potential impact of the wait times (Sections 5.2.4, EC.5.2, and EC.5.3) of our policy on outcomes is to examine changes in an individual's circumstances over time under no treatment, which we can observe from retakes of the VI-SPDAT. Because case managers typically administer a retake of the VI-SPDAT only when they notice a visible shift in circumstances, the frequency of retakes, the time elapsed between them, and the magnitude of any score change together provide a proxy for how quickly vulnerability actually worsens. For example, if individuals rarely retake the assessment within 100 days (the 90<sup>th</sup> quantile of wait time for PSH under our policy) and their scores do not increase significantly (indicating worsening vulnerability), then this suggests wait times at the levels of our policy do not result in material changes to circumstances.

To analyze individual changes based on VI-SPDAT retakes, we subset our data on those who received no treatment historically since these individuals represent what happens to someone as they wait for their allocation. We see that only 6.7% of all individuals had a reassessment within 100 days. Among those who ever retook the assessment, the average and median time to retake were 385 and 303 days, respectively. These numbers indicate that case workers did not find most



**Figure EC.14** DoubleML/R-Learner causal estimate of the heterogeneous negative impact of increasing wait times on probabilities of positive outcomes.

individuals' circumstances to change materially within the time frame of the wait times of our policy to administer a reassessment. Furthermore, amongst those who needed a reassessment within a year, only 51% of the individuals had a score increase (their circumstances worsened) and the mean and median score changes both increased by 1 on an 18 point scale. Consequently, the wait times our policy introduces are short relative to the timescales on which VI-SPDAT scores (a proxy for individuals' situations) tend to change, suggesting the practical acceptability of our policy's queuing wait times.