LEARNING TREATMENT EFFECTS DURING RESOURCE ALLOCATION VIA PRIORITY-QUEUE RANDOMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Public service programs often allocate limited resources under uncertainty about their benefits, creating a need for randomization to support credible evaluation. In practice, however, applicants commonly enter waitlists where resources are prioritized toward individuals judged to have higher need through tiered priority queues, making direct randomization difficult. Motivated by this, we develop an experimental design framework for learning treatment effects while treating those most in need where incoming applicants are randomized into priority queues based on their assessed risk scores. Treatments are then provided across queues in priority order and first-in-first-out within queue as budget becomes available. Our contributions are two-fold. First, we characterize what causal effects are identified under this priority-queue allocation. When arrivals are exogenous, treatments are conditionally randomized, and hence standard estimands are identified; when arrivals are endogenous, queue randomization instead provides an instrument for treatment, identifying local treatment effects induced by the queuing process. Second, we develop optimized queue-assignment designs that trade off statistical efficiency against prioritizing higher-need applicants. We show in the process that, despite dependence in treatment assignments induced by the design, usual iid efficiency bounds remain well-justified design objectives. We illustrate our designs using data from a housing allocation program in a large U.S. county.

1 INTRODUCTION

In many public service programs, scarcity is managed not through a one-time allocation decision, but through tiered priority queues. Applicants arrive over time, resources become available over time, and each newly available slot is offered according to a service rule: applicants in higher-priority tiers are served before those in lower-priority tiers, while applicants within the same tier are typically ordered by time spent waiting. Such priority-based waiting systems are widely used, including for child care subsidies (Virginia Department of Education, 2024; 2025), Medicaid home- and community-based services (Medicaid and CHIP Payment and Access Commission (MACPAC), 2020), public housing and homelessness assistance (Allegheny County Department of Human Services, 2020), and related settings such as organ transplantation (Colvin-Adams et al., 2012).

While these programs provide essential services, they often operate under significant uncertainty regarding their causal impact. Rigorous randomized evaluation is the gold standard for resolving this uncertainty, yet direct randomization, such as a lottery for treatment slots, is often administratively or ethically unfeasible in these contexts. Bypassing high-priority applicants to serve those with lower assessed need would violate the program's core mandate.

We therefore propose an experimental design framework that resolves this tension by randomizing queue assignment rather than treatment allocation. The resulting design problem is substantially more complex because treatment is not assigned independently at a fixed point in time; applicants arrive sequentially, resources become available over time, and each treatment decision depends on the evolving composition of all priority queues. Moreover, the design must preserve the program's priority-based logic. Under our design, applicants are randomized into priority queues with probabilities depending on their risk scores. Once assigned, treatment assignment is governed by existing rules: resources are allocated across queues in priority order and first-in-first-out within each queue.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Our contributions are two-fold. First, we characterize what causal effects are identified under this dynamic allocation process. We show that when arrivals are exogenous, treatment is conditionally randomized and standard treatment effects can be recovered; when arrivals are endogenous, queue randomization instead provides an instrument for treatment, identifying local effects induced by the queuing process. Second, we develop optimized queue-assignment designs that balance statistical efficiency with the prioritization of high-need applicants. This design problem is difficult because the policymaker controls queue-assignment probabilities rather than treatment propensities directly, and because treatment assignments are dependent across applicants through the queuing process. We formulate policies that maximize statistical efficiency subject to a utility constraint ensuring that high-need individuals remain likely to be prioritized.

Relevant literature: Our paper builds on Wilder & Welle (2025) who study how to design randomized allocation rules that balance two competing goals: prioritizing individuals with greater need and learning treatment effects. We extend this perspective to settings where treatment is instead allocated through priority-based rules. Our work is also related to the market-design literature on dynamic waitlists and priority-based allocation, studies settings in which agents and/or objects arrive over time, and evaluates mechanisms according to objectives such as welfare, fairness, and incentives to accept or defer offers e.g., Thakral (2019); Arnosti & Shi (2020); Schummer (2021); Leshno (2022); Murra-Anton & Thakral (2026). Our focus is not on characterizing the welfare or incentive properties of a waitlist mechanism per se, but on designing randomization within such mechanisms to enable causal inference while preserving priority-based allocation goals.

Our design criterion is related to the broader literature on adaptive experimental design for efficient treatment effect estimation that uses covariates, pilot data, or sequential information to improve precision (Hahn et al., 2011; Tabord-Meehan, 2023; Kato et al., 2025). Related work extends this logic to indirect experiments and settings with noncompliance (Chandak et al., 2024; Oprescu et al., 2026). A separate literature studies experimental design and policy learning under capacity constraints (Yamin et al., 2025; Zhou, 2023; Tang et al., 2023; Sun et al., 2021). These papers share our interest in scarce resource allocation, but typically choose allocation policies directly, whereas in our setting treatment is induced by a subsequent queuing process. Closest to our inferential setting is perhaps De Chaisemartin & Behaghel (2020), who study causal estimation under randomized waitlists induced by limited resources. There, applicants' ranks are randomized and offers proceed down the list until seats are filled, which conflicts with the priority structure we seek to preserve.

2 Basic setup and queue-randomizing design

We consider a setting in which n units arrive into our system over a fixed enrollment horizon. Each unit has an arrival time $A \in [0, \tau]$ and observed covariates $X \in \mathbb{R}^p$. Treatments are distributed at discrete allocation periods $t \in \{1, \ldots, \tau\}$, where $\tau \in \mathbb{N}$ is fixed. We use $Z \in \{0, 1\}$ to denote terminal treatment status, where $Z = 1$ means that the unit is treated at some point before the end of the allocation horizon and $Z = 0$ otherwise. We adopt the potential outcomes framework for causal inference (Neyman, 1990; Rubin, 1974) and write $Y^0, Y^1 \in \mathbb{R}$ as the (stationary) potential outcomes under no treatment and treatment, respectively. We therefore posit that we have n iid draws $O = (A, X, Y^0, Y^1)$ from some joint distribution $\mathcal{P}$. We make the standard *consistency* assumption throughout that the observed Y is $Y = Y^z$ if $Z = z$.

Unlike idealized settings where treatment is determined independently for each individual, many public service settings allocate scarce resources through tiered priority queues. There are K queues, indexed by $\{1, \ldots, K\}$, where lower-indexed queues have higher priority. At each individual's arrival time A, they are assigned a queue $Q \in \{1, \ldots, K\}$ based on their observable covariates X. At the end of each allocation period, a budget b becomes available, representing b available treatment slots, such as housing vacancies. In the formal notation below, we write this fixed period-specific budget as $b_{t,n}$ to allow treatment capacity to vary across periods and with the system size n, and $b_n \equiv (b_1, \ldots, b_{\tau,n})$ for the entire sequence of budgets. At the end of period t, the policymaker considers all units who have arrived by that period and have not yet been treated. Treatment is offered first to individuals in the top-priority queue 1, then to individuals in queue 2, and so on, until the period-t budget $b_{t,n}$ is exhausted. Within each queue, treatment is first-in, first-out (i.e., treatment is offered in order of arrival time). Units who are not treated in period t remain eligible for treatment in later periods, which therefore yields a waitlist structure.

We study a class allocation policies where the policymaker chooses which queue an individual with covariates X is assigned to via a (potentially randomized) policy $\theta(X)$, where $\theta(X) = (\theta_1(X), \ldots, \theta_K(X))^\top \in \Delta_K$ denotes the vector of queue assignment probabilities from which the queue $Q \in \{1, \ldots, K\}$ is drawn (throughout, $\Delta_K \equiv \{w \in \mathbb{R}_+^K : \sum_{k=1}^K w_k = 1\}$ denotes the simplex). Since θ depends in X, the policymaker can offer different assignment probabilities depending on individual characteristics. At one extreme, higher-risk individuals might be deterministically assigned to higher-priority queues, as is typically the case in practice. At the other end, θ could be uniform for all individuals (effectively implementing an RCT). In between, higher risk individuals can be offered a greater probability of placement in top-priority queues. The choice of queue for an individual induces the probability that they will actually receive treatment. Let $\mathbb{P}_{\theta,b_n}$ be the distribution of the treatment induced by our design. We use $\tilde{\pi}_n(x, q; \theta, b_n) \equiv \mathbb{P}_{\theta,b_n}(Z = 1 \mid X = x, Q = q)$ to denote the *queue-conditional propensity score* and $\pi_n(x; \theta, b_n) \equiv \mathbb{P}_{\theta,b_n}(Z = 1 \mid X = x)$ to denote the *marginal propensity score* conditional just on covariates X.

Queue-based structure is needed to manage the operational complexity of real-world services, where applicants and resources both arrive online. However, it creates substantial complexity for inference of treatment effects because whether an individual is offered treatment depends on how the arrival of other individuals and resources flow through the queuing system. This stands in stark contrast to many other allocative settings, where a simple cutoff for eligibility enables designs like regression discontinuities. Our goal is two-fold. First is to determine *which causal effects are identified under this allocation mechanism* given the complexity of the process. Second is develop experimental designs – *how can a policymaker select a randomized policy θ* enable accurate estimation of treatment effects while still allocating treatment to those most in need?

3 What estimands can be identified?

The most basic question is what causal effects can be identified under what conditions for queue-based allocation systems. So far, we have not imposed any assumptions on how the arrival times A relate to the other variables. This turns out to be a critical distinction because the arrival time distribution determines how resources are prioritized between otherwise-similar individuals.

The simplest setting would be if the arrival time were *independent* of the potential outcomes, at least conditional on observables: $A \perp\!\!\!\perp (Y^0, Y^1) \mid X$. In this setting, arrival times create a "natural experiment" of their own, apart from any explicit randomization performed by the policymaker. That is, individuals who arrive when there happens to be more budget available may be more likely to receive treatment than otherwise-similar individuals who arrive when demand is higher and the system more congested. Formalizing this intuition, we show:

Theorem 1. *Suppose that the arrivals satisfy $A \perp\!\!\!\perp (Y^0, Y^1) \mid X$. Then, treatment assignments are conditionally randomized: $Z \perp\!\!\!\perp (Y^0, Y^1) \mid X$. Moreover, suppose that every level of the covariates X satisfies positivity: there exist $\gamma > 0$ and $n_0 \in \mathbb{N}$ such that for every $n \geq n_0$, $\gamma \leq \pi_n(X; \theta, b_n) \leq 1-\gamma$. Then, any functional of the potential outcome distribution $\mathbb{P}(Y^z \mid X = x)$ is identified. For example, the average treatment effect (ATE) $\psi_{\text{ATE}} \equiv \mathbb{E}(Y^1 - Y^0)$ is identified.*

This demonstrates that assuming exogeneity of arrival times suffices to identify many causal effects of interest simply by comparing individuals who are vs are not treated (conditional on covariates). Such comparisons have been used in a number of applied program evaluation settings, and Theorem 1 helps clarify the assumptions on the arrival process necessary for this strategy to be valid (which have not been discussed in previous resources for practitioners, e.g. Frechtling (2010)).

Making these assumptions explicit reveals that they are potentially stringent when arrivals may be correlated with underlying need. For example, consider a public housing program. It could be the case that more people apply for public housing assistance in time periods where outcomes are worse (e.g., because of worse overall economic circumstances), in ways that may not be fully captured by the covariates. In such a case, the naive analysis suggested by Theorem 1 would suffer from confounding. Hence, we now turn to the development of identification strategies which do not rely on the assumption $A \perp\!\!\!\perp (Y^0, Y^1) \mid X$.

Discarding the assumptions of exogenous arrivals substantially complicates the identification of treatment effects. Even if the policymaker explicitly randomizes which queues individuals are placed into via the choice of θ, this does not suffice for treatment assignments to be randomized: endogeneity

in the arrival process can still confound treatment assignment by making individuals who arrive at particular times (and have particular potential outcomes) be more likely to receive treatment.

To resolve these difficulties, we propose an identification strategy which treats the random assignment of the queue as an instrumental variable (IV). Since the queue assignment can be explicitly randomized by the policymaker, standard IV assumptions will be satisfied by design.[1] Following the literature on formula instruments (Borusyak et al., 2025), we propose to instrument Z using the difference between a given individual's *expected* treatment probability pre-randomization and their *realized* treatment probability post-randomization:

$$r_n(X, Q) \equiv \tilde{\pi}_n(X, Q; \theta, b_n) - \pi_n(X; \theta, b_n).$$

r_n accounts for the fact that individuals with particular covariates X might be more likely ex-ante to be placed into queues with high treatment probabilities by recentering relative to $\pi_n(X; \theta, b_n)$.

Lemma 1. *Under our design, $r_n(X, Q)$ satisfies conditional IV-unconfoundedness, exclusion, and the induced treatment statuses satisfy monotonicity.*

A natural causal estimand is therefore the IV ratio:

$$\beta_{\text{IV}} \equiv \frac{\mathbb{E}\{r_n(X, Q)Y\}}{\mathbb{E}\{r_n(X, Q)Z\}} = \frac{\mathbb{E}\big[\{\tilde{\pi}_n(X, Q; \theta, b_n) - \pi_n(X; \theta, b_n)\}Y\big]}{\mathbb{E}\big[\{\tilde{\pi}_n(X, Q; \theta, b_n) - \pi_n(X; \theta, b_n)\}Z\big]},$$

given the relevance assumption $\mathbb{E}\{r_n(X, Q)Z\} \neq 0$. However, it is not apriori clear exactly what treatment effect β_{IV} targets. Generally, IVs are interpretable as identifying a treatment effect on a latent complier population who are treated only "because" of the instrument. Here, this complier population is induced by the queuing process. We next derive an explicit interpretation of β_{IV} as a particular weighted combination of local average treatment effects (LATEs).

Formally, for any two queues $k < \ell$, let $\mathcal{C}_{k,\ell}$ denote the set of individuals who (because of their arrival times) would be treated if randomized to queue k but *not* if randomized to queue ℓ. It turns out β_{IV} corresponds to a weighted average of treatment effects for such latent complier populations:

Theorem 2. *Let $\psi_{k,\ell}(X) \equiv \mathbb{E}(Y^1 - Y^0 | X, \mathcal{C}_{k,\ell})$ and $w_{k,\ell}(X) \equiv \theta_k(X)\theta_\ell(X)(\tilde{\pi}_n(X, k; \theta, b_n) - \tilde{\pi}_n(X, \ell; \theta, b_n))^2$. Then,*

$$\beta_{\text{IV}} = \frac{\mathbb{E}\{\sum_{k<\ell} w_{k,\ell}(X)\psi_{k,\ell}(X)\}}{\mathbb{E}\{\sum_{k<\ell} w_{k,\ell}(X)\}}.$$

The weights $w_{k,\ell}(X)$ reflect that particular LATEs $\psi_{k,\ell}(X)$ are weighted more heavily if there is a larger probability of individuals with covariates X being randomized to queues k and ℓ and if queues k and ℓ induce substantially different treatment probabilities. That is, such designs target individuals whose treatment status is more likely to change because of randomization.

In practice, β_{IV} might be hard to interpret if there is little basis for identifying which individuals are treated only at specific levels of priority. To provide a simplified estimand for interpretation, and later for experimental design, we also study the partially linear IV model

$$Y = \psi_Z Z + g(X) + U, \qquad \mathbb{E}(U \mid X, Q) = 0, \tag{1}$$

studied by, e.g., Robinson (1988) and Chernozhukov et al. (2018). This model estimates a single homogeneous treatment effect ψ_Z, although the covariates are still allowed to have a separate (nonparametric) effect on the outcome via $g(X)$. For practitioners interested in estimating such an "average" effect, we show that a scaled version of r_n remains the uniquely best choice of instrument for attaining the usual iid efficiency bound (Chamberlain, 1992):

Theorem 3. *Let $\sigma(X) \equiv \mathbb{E}(U^2|X)$. Then, $r_n(X, Q)/\sigma(X)$, up to multiplicative scaling, uniquely attains the iid efficiency bound for the partially linear IV model equation 1.*

Intuitively, $r_n(X, Q)$ captures the randomized first-stage variation generated by queue assignment after partialing out X. Scaling by $1/\sigma(X)$ then prioritizes covariate regions with less residual outcome noise, yielding the usual inverse-variance weighting logic.

[1] Exclusion is justified here because queue assignment is an administrative priority label whose only operational role is to determine the order in which applicants are offered treatment, and hence whether they receive treatment before the end of the allocation horizon. Although queue assignment may affect waiting time, our stationary potential outcome formulation assumes that outcomes depend only on terminal treatment receipt, not on the exact timing. This restriction would fail if queue assignment changed other aspects of service delivery or if outcomes depended directly on treatment timing.

4 How should experiments be designed?

So far, we have treated the queue assignment policy $\theta(X)$ as fixed. Next, we turn to the *design* of experiments that use randomization to queues. In particular, we consider policymakers who wish to balance learning treatment effects against the loss that randomization creates with respect to the originally desired prioritization. Following Wilder & Welle (2025), we formalize this by modeling the policymaker as having a utility $u(X)$ for offering treatment to a participant with covariates X. For example, u could be a risk score, representing a preference for allocating treatment to higher-risk participants. If the policymaker has no desire for experimentation, they would simply assign queues deterministically based on u, e.g., assigning the highest-risk applicants to the highest priority queue, and so on. However, without randomization, it becomes difficult to credibly infer treatment effects (particularly without assuming exogeneity of arrivals). On the other hand, uniformly random assignments would maximize power to estimate treatment effects but entirely ignore u.

Our goal is to design policies θ which lie on the Pareto curve between these extremes, allowing policymakers to optimally trade off between the two goals. Wilder & Welle (2025) showed that in the case of single-shot treatment decisions, this can be reduced to a convex optimization problem that trades off expected utility with the optimal variance (efficiency bound) for estimating the treatment effect. However, the dynamic setting with queue-based allocation creates substantial new challenges. First, we optimize over the queue allocation probabilities θ instead of directly over the treatment propensities π; formulating an optimization problem requires a tractable structure for the map $\theta \mapsto \pi$. Second, in our setting, treatment assignments are not independent across individuals: the queue process creates correlations since individuals "compete" for fixed resources. Without iid data, it is unclear whether standard efficiency bounds apply. Third, particularly under endogeneous arrivals, the treatment effects that we can estimate are substantially more complicated than a simple average treatment effect and it is not clear how to derive a tractable variance proxy for optimization.

To address the first and second challenge, we describe a set of assumptions on the queuing process under which we can formulate tractable proxy objectives for optimization. We emphasize that these assumptions are *not* required for valid identification of treatment effects: our results in the previous section hold regardless. Instead, we view them as a useful approximation that can be imposed at the design stage in order to guide the planning of an experiment. All asymptotic statements are taken with respect to a growing experiment size $n \to \infty$, holding fixed the number of queues K, the number of allocation periods τ, the queue-assignment policy θ, and the arrival-time distribution.

Assumption 1 (Fixed queue proportions)**.** $\mathbb{E}\{\theta(X)\} = p \equiv (p_1, \ldots, p_K)^\top \in \Delta_K$.

Assumption 2 (Stable and non-negligible budget)**.** *There exists* $\beta \in (0, 1)$ *such that, as* $n \to \infty$*,*

$$\max_{1 \le t \le \tau} \left| \frac{1}{n} \sum_{s \le t} b_{s,n} - \beta \mathbb{P}(A \le t) \right| \to 0.$$

Assumption 3 (Independent arrivals)**.** *The arrivals are independent of covariates, i.e.,* $A \perp\!\!\!\perp X$.

We view Assumption 1 as a mild condition that is enforceable by-design: it says that each queue should capture some fixed proportion of applicants in-expectation. This is true of many empirical examples, where, e.g., a queueing system with 3 levels is designed so that the top queue corresponds to the 1/3 highest-risk individuals. Assumption 2 says that the budget is released over time in proportion to enrollment and that, asymptotically, treatment capacity is a constant fraction of the experiment size. If the fraction treated vanishes, then there is not an appropriate asymptotic sense in which the amount of the data observed increases with n.

Assumption 3 is perhaps the most substantive, and says that arrival times do not sort systematically according to covariates. This can be viewed as optimizing under a stationarity condition, where the arrival distribution does not change significantly over time. We view this as a reasonable approximation for the design stage since the policymaker likely does not have the ability to predict with any precision how the future arrival distribution will change. Additionally, without such an assumption, the relevant asymptotics are hard to define since the queuing system and corresponding treatment propensities need not converge as the number of individuals grows.

Correspondence between queues and treatment propensities: Under these conditions, we address the first challenge by showing that the policy θ induces a tractable mapping to the resulting

treatment propensities π. The relevant object for later analysis is the *asymptotic propensity score*

$$\pi(x;\theta) \equiv \lim_{n\to\infty} \pi_n(x;\theta,b_n) = \lim_{n\to\infty} \sum_{k=1}^{K} \tilde{\pi}_n(x;\theta,b_n)\theta_k(x). \tag{2}$$

We show $\pi(x;\theta)$ takes a specific form:

Theorem 4 (Design-induced asymptotic propensity score)**.** *Suppose Assumptions 1–3 hold. Let $c_k \equiv \sum_{j=1}^{k} p_j$ with $c_0 \equiv 0$. For each k, define*

$$\alpha_k(\beta,p) \equiv \begin{cases} \{(\beta - c_{k-1})_+ - (\beta - c_k)_+\}/p_k, & p_k > 0, \\ 0, & p_k = 0, \end{cases}$$

where $u_+ \equiv \max\{u, 0\}$. Then, for a.e. x, $\pi(x;\theta) = \sum_{k=1}^{K} \alpha_k(\beta,p)\theta_k(x)$.

Intuitively, this corresponds to modeling individuals as having a fixed probability of treatment conditional on their queue assignment, with the final treatment propensity being a mixture of these queue-specific probabilities. Thus the asymptotic propensity score is affine in θ, which is the key structural property that makes the design problem tractable.

Estimation with non-iid assignments: We address the second challenge by showing that, even with dependent treatment allocations, estimators based on the usual iid efficient influence function (EIF) or orthogonal estimating equation representations (Kennedy, 2024) have the same leading asymptotic variance as in the iid case. We start with the case of exogenous arrivals, focusing on the standard ATE estimand. Let $\mu_z(x) \equiv \mathbb{E}(Y \mid Z = z, X = x)$ and $\mathbb{P}_n f \equiv n^{-1}\sum_{i=1}^{n} f(O_i)$ denote the empirical average. A canonical EIF-based estimator is the doubly-robust estimator

$$\widehat{\psi}^{\mathrm{DR}}_{\mathrm{ATE}} \equiv \mathbb{P}_n\left[\widehat{\mu}_1(X) - \widehat{\mu}_0(X) + \frac{Z}{\pi_n(X;\theta,b_n)}\{Y - \widehat{\mu}_1(X)\} - \frac{1-Z}{1-\pi_n(X;\theta,b_n)}\{Y - \widehat{\mu}_0(X)\}\right].$$

Its asymptotic variance under our treatment allocation scheme can be characterized as follows:

Theorem 5 (Asymptotic variance of $\widehat{\psi}^{\mathrm{DR}}_{\mathrm{ATE}}$)**.** *Suppose the conditions in Theorem 1 and Assumptions 1–3 hold. Further, assume $\widehat{\mu}_z$ is constructed from an independent sample $\mathcal{D}$ such that $\|\widehat{\mu}_z - \mu_z\|_{L_2(\mathbb{P})} = o_{\mathbb{P}}(1)$ for $z \in \{0,1\}$. If there exists $m < \infty$ such that $\mathbb{E}\{(Y^1)^2 + (Y^1)^0\} \le m$, then conditional on $\mathcal{D}$,*

$$\lim_{n\to\infty} n\mathrm{Var}(\widehat{\psi}^{\mathrm{DR}}_{\mathrm{ATE}}) = \mathbb{E}\left\{\frac{\mathrm{Var}(Y \mid Z=1, X)}{\pi(X;\theta)} + \frac{\mathrm{Var}(Y \mid Z=0, X)}{1-\pi(X;\theta)}\right\} + \mathrm{Var}\big(\mu_1(X) - \mu_0(X)\big).$$

This result has two implications. First, it characterizes when we can analyze the data from the queue-based experiment *as if* the treatment assignments were iid and still obtain valid confidence intervals based on asymptotic normality (Appendix B). Second, the variance above is exactly the semiparametric efficiency bound for estimating the ATE from iid data (Hahn, 1998; Imbens & Wooldridge, 2009). Thus, the same quantity gives the accuracy with which treatment effects can be estimated even when treatment assignments are correlated via the queuing process. This will allow us to use the efficiency bound to instantiate an objective for optimizing over experimental designs.

The key intuition behind Theorem 5 is that the dependence induced queuing systems, though complex, is structured. There are two sources of randomness: the queue composition and the arrival order within partially served queue-period cells. Under our assumptions, queue composition is stable in the sense that changing one other unit's queue changes each queue count by at most one, and treated queue counts by at most a constant. Conditional on the realized design variables, treatment is deterministic outside partially served cells, while treatment within such cells behaves like sampling without replacement from a cell of size order n. Thus pairwise treatment dependence is at most order n^{-1}. Importantly, this dependence does not affect the leading oracle term: after decomposing the DR estimator into an average of oracle influence functions plus a nuisance-estimation remainder, the dependence enters only through the nuisance remainder, where it is multiplied by nuisance estimation errors. Because $\widehat{\mu}_z$ are consistent, the remainder contributes $o_{\mathbb{P}}(1)$ to the scaled variance, leaving the usual influence-function variance.

Estimation with endogenous arrivals: we now show an analogous result holds for the case of endogenous arrivals, when randomization at the queue level provides an instrumental variable for the treatment. As discussed above, the definition and interpretation of treatment effects is more complex in this setting. The fully nonparametric characterization of the IV estimand show that it is a weighted combination of LATEs where the weights depend themselves on the choice of instrument (i.e., the policy θ we are now optimizing over). To provide a tractable objective for the design stage, we suggest optimizing with respect to estimation efficiency for the treatment effect in the partially linear IV (PLIV) model, which provides a natural counterpart to estimating a homogeneous average treatment effect. We show an analogous result to Theorem 5 holds for an estimator obtained by solving an empirical moment equation induced by the instrument in Theorem 3:

$$\widehat{\psi}_Z^\star \equiv \frac{\mathbb{P}_n\left[\frac{r_n(X,Q)}{\widehat{\sigma}(X)}\{Y-\widehat{m}(X)\}\right]}{\mathbb{P}_n\left[\frac{r_n(X,Q)}{\widehat{\sigma}(X)}\{Z-\pi_n(X;\theta,b_n)\}\right]}, \qquad m(x)\equiv\mathbb{E}(Y\mid X=x).$$

Theorem 6 shows that, under our design, the asymptotic variance of $\widehat{\psi}_Z$ matches the usual efficiency bound expression evaluated at the limiting instrument $\alpha_Q(\beta,p)-\pi(X;\theta)$ (Chamberlain, 1992).

Theorem 6 (Asymptotic variance of $\widehat{\psi}_Z^\star$)**.** *Suppose Assumptions 1–3 and the model equation 1 hold. Further, assume $\widehat{\sigma}$, $\widehat{m}$ are constructed from an independent sample $\mathcal{D}$ such that $\|\widehat{\sigma}-\sigma\|_{L_2(\mathbb{P})}=o_{\mathbb{P}}(1)$ and $\|\widehat{m}-m\|_{L_2(\mathbb{P})}=o_{\mathbb{P}}(1)$. If $\mathbb{E}\left[\{\alpha_Q(\beta,p)-\pi(X;\theta)\}^2\right]\geq\gamma>0$, then conditional on $\mathcal{D}$,*

$$\lim_{n\to\infty} n\mathrm{Var}(\widehat{\psi}_Z^\star)=\left(\mathbb{E}\left[\frac{\{\alpha_Q(\beta,p)-\pi(X;\theta)\}^2}{\sigma(X)}\right]\right)^{-1}.$$

Multiobjective optimization: In either setting (exogenous or endogenous arrivals), our goal now is to select the queue assignment probabilities θ to optimally trade off between the participant's utility $u(X)$ and statistical power to estimate treatment effects. We use $\mathbb{V}_\psi(\theta,\eta)$ to denote the asymptotic variance of an estimate of treatment effect ψ given queue assignment probabilities θ and nuisance functions η. η comprises terms in the objective that are not knowable ahead of running the experiment. For example, for the ATE estimand above, $\eta=(\mathrm{Var}(Y|Z=1,\cdot),\mathrm{Var}(Y|Z=0,\cdot))$ since the conditional variances of the outcomes typically cannot be estimated precisely before running an experiment. Similarly, for the PLIV case, $\eta=\sigma(\cdot)$. We stipulate that such nuisances lie within an uncertainty set $\mathcal{U}$ and solve a robust optimization problem of the form

$$\min_{\theta:\mathcal{X}\to\Delta_k}\max_{\eta\in\mathcal{U}}\mathbb{V}_\psi(\theta,\eta)\text{ s.t. }\mathbb{E}[\theta(X)]=p;\ \mathbb{E}[\pi(X;\theta)u(X)]\geq c,$$

where the first requires that the queue assignment probabilities θ satisfy the budget constraint, and the second imposes the expected utility of the resulting treatment assignments be at least c. By adjusting the value of c, we can trace out the entire Pareto curve between treatment effect variance and expected utilities, allowing the policymaker to select a point that represents their desired balance. Note that since π is affine in θ, all of the constraints are linear with respect to θ.

As a general approach to constructing the uncertainty set $\mathcal{U}$, we assume that the nuisances are bounded, i.e., $\mathrm{Var}(Y|Z,X)\leq\delta$ for some constant δ for the ATE or $\sigma(X)\leq\delta$ for the PLIV model. Generically, we represent this as an uncertainty set of the form $\mathcal{U}=\{\eta:||\eta||_\infty\leq\delta\}$. If the policymaker happens to have additional information (e.g., a tighter bound on one of the conditional variance functions), our techniques can be naturally extended. For uncertainty sets that have such box-constraint style structure, we obtain tractable simplifications of the inner max:

Theorem 7. *For the ATE case, minimizing $\max_{\eta\in\mathcal{U}}\mathbb{V}_{\mathrm{ATE}}(\theta,\eta)$ is equivalent to minimizing $\mathbb{E}\left\{\frac{1}{\pi(X;\theta)}+\frac{1}{1-\pi(X;\theta)}\right\}$, which is convex in θ. For the PLIV case, minimizing $\max_{\eta\in\mathcal{U}}\mathbb{V}_{\mathrm{PLIV}}(\theta,\eta)$ is equivalent to maximizing $\mathbb{E}\left[\{\alpha_Q(\beta,p)-\pi(X;\theta)\}^2\right]$, which is concave in θ.*

One final complication is that since π is affine in θ, there may not be a unique solution to the above optimization problems. We resolve this by augmenting the objective with a regularizer $\kappa R(\theta)$ where R is strictly convex and $\kappa>0$. Experimentally, we take R to be the (negative) entropy to tie-break in favor of designs with greater randomization at the queue level. The final result is a strictly convex minimization (or strictly concave maximization) with respect to θ.

To compute the resulting policies in practice, we work with the dual formulation of the constrained population problem that yields a tractable finite-sample analogue. We defer the explicit sample optimization problem and additional details to Appendix A.

5 NUMERICAL EXPERIMENTS

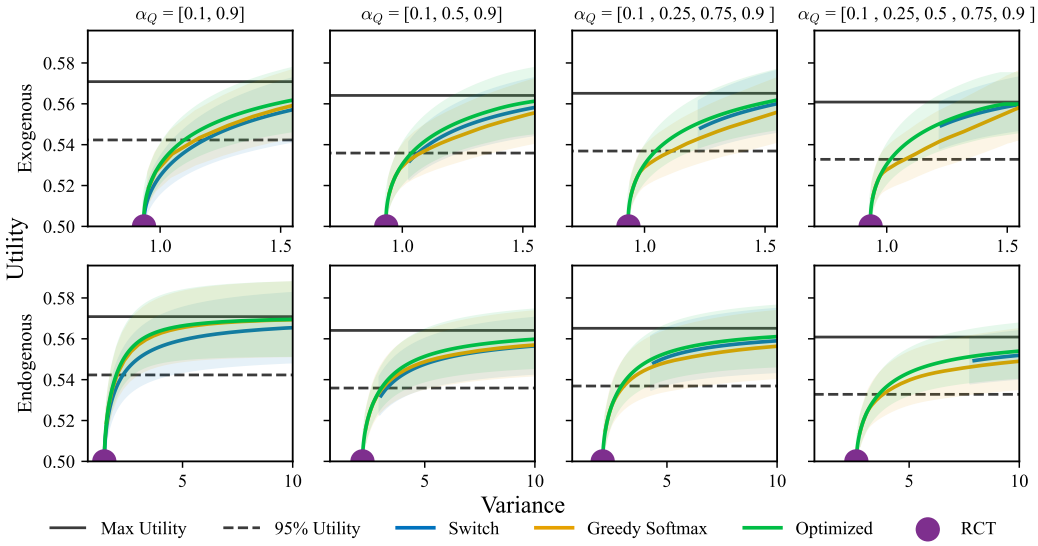

Figure 1: Pareto frontier of variance and utility. under our optimized designs (green) versus two heuristic designs (orange, blue) and an RCT baseline (purple) under exogenous arrivals (top) and endogenous arrivals (bottom) across different queue configurations with a 50% treatment budget. Queues are equally sized but vary in number, with treatment probabilities at the top of each column. The optimized designs achieve higher efficiency and utility across all configurations compared to the heuristic methods. Generally, exogenous designs are more efficient than endogenous designs.

We evaluate the sample efficiency and robustness to arrival endogeneity of our proposed designs using data from a public housing allocation program in a large U.S. county. The program assigns individuals to priority queues using a predictive risk score of homelessness, $h(X)$, built from administrative records, and allocates housing units using exactly the kind of first-in, first-out assignment mechanism studied here. The dataset contains 2,317 service-eligible individuals and spans the last 10 years. Our outcome of interest is an indicator of whether an individual experiences homelessness in the six months following unit assignment. Throughout, we set utility $u(X) = h(X)$ to reflect a preference for assigning units to those with the highest predicted risk of homelessness. As only outcomes under the assigned treatment are observed, we use the predictive risk scores to simulate potential outcomes with a homogeneous treatment effect. We provide further details on the data and synthetic data-generating process in Appendix D and an experiment using publicly available data in Appendix E.

5.1 SAMPLE EFFICIENCY

To evaluate the efficiency of our proposed estimators, we consider equally-sized queues under a 50% treatment budget across various queue treatment probabilities. We include results under alternative queue configurations and budgets in Appendix D.3. We then solve the optimization problem across a range of utility constraints for both the exogenous and endogenous arrivals designs. Figure 1 shows the Pareto frontier between utility and asymptotic variances in Theorems 5 and 6. We compute 95% confidence bands using the multiplier bootstrap with 10,000 replicates (Kennedy, 2019; Van Der Vaart & Wellner, 1996). We compare our proposed designs to two heuristics for constructing queue assignment probabilities: (1) *Switch* assigns individuals to queues deterministically via $u(X)$, then randomizes each assignment via switching probabilities to adjacent queues weighted by relative queue size, (2) *Greedy Softmax* sequentially constructs assignment probabilities from highest- to lowest-priority queues using a softmax-style score over remaining utility. In order to enforce queue proportions, Greedy Softmax requires access to the full sample. Both methods include a parameter governing assignment determinism, which we vary to compute the Pareto curve. Additional implementation details are provided Appendix D.3.

Figure 1 reveals a favorable concave trade-off between utility and statistical efficiency (variance) under our proposed designs. Both designs achieve high utility at modest cost to variance relative to an RCT. The endogenous design is generally higher variance than the exogenous design, and becomes increasingly so as the number of queues increases. Our optimization approach generally outperforms both heuristics, achieving consistently high statistical efficiency and utility across all queue configurations. The Switch heuristic's restriction to neighboring queues means that it is only capable of generating a portion of the full tradeoff curve, particularly when the number of queues is large. For a smaller number of queues, its performance is dominated by the optimized design. On the other hand, the performance of Greedy Softmax deteriorates as more queues are added. As the number of queues increases, more individuals fall near queue boundaries, and the greedy allocation spreads their residual mass across an increasing number of lower queues rather than concentrating mass in the adjacent queue. The optimized design avoids both failure modes by jointly solving for assignment probabilities across all queues simultaneously.

5.2 Robustness to Arrival Endogeneity

Although generally more sample efficient, the design under exogenous arrivals may lead to biased estimates if arrival times are not conditionally independent of potential outcomes. To demonstrate, we introduce a synthetic confounder U to induce arrival endogeneity. We then estimate $\widehat{\psi}_Z$ and $\widehat{\psi}^{\text{DR}}$ using the known synthetic values and calculate the bias of each estimator, see Appendix D.4 for additional implementation details. Figure 2 shows the average bias of both designs across utility levels for two equally-sized queues under a 50% budget and varying treatment probabilities. For each queue configuration, we solve the optimization problem to obtain queue assignment policies under both designs across a range of utility constraints, and compute bias over 10,000 simulations. As expected, the endogenous estimator (dotted) exhibits little bias across the utility levels considered. The exogenous estimator (dashed) exhibits substantial bias at high levels of utility and low treatment probabilities in the highest-priority queue. As the treatment probability of the highest-priority queue α_1 increases, treatment becomes mostly determined by queue assignment rather than arrival order, thus mitigating arrival endogeneity.

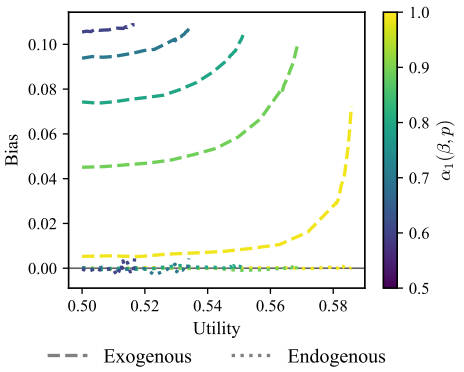

Figure 2: Bias under arrival endogeneity for two queues with varying treatment probabilities. The endogenous estimator (dotted) exhibits little bias; the exogenous design (dashed) shows substantial bias at high utility levels and when treatment probability in the highest-priority queue is low.

6 Conclusion

In this paper, we introduced an experimental design framework that integrates randomized queue assignment with the priority-based allocation rules that are standard in public service settings. We first characterized causal identification under two distinct arrival regimes and developed optimized queue-assignment policies that effectively balance statistical efficiency against the utility of prioritizing high-need applicants. Our current work focuses on whether treatment is received at all before the allocation horizon ends. While this is natural when the primary scientific question concerns overall access to treatment, the timing of the intervention may itself be of substantive importance. In such cases, it would be valuable to extend our framework to accommodate time-indexed treatment probabilities or the full distribution of treatment timing.

Overall, our results show that incorporating randomization into priority-queue allocation schemes offers policymakers a means to learn about program effectiveness while limiting sacrifices in the targeting of the program and accommodating the operational constraints of dynamic settings. While the right balance between objectives is setting-specific, our hope is that expanding the Pareto frontier opens up the possibility for more widespread evaluation and improvement of critical services.

AI USE STATEMENT

In this work, we used generative AI tools for generating synthetic datasets, assisting in the writing of proofs, designing and providing feedback on experiments, implementing methods, and cleaning and reformating dataset. We have not used generative AI tools for helping develop theoretical models or conceptual frameworks, formulating mathematical claims, providing critical ingredients for proving mathematical claims, proposing or refining hypotheses, designing or providing feedback on research methodology or experiments, supporting qualitative and thematic data analysis, and interpreting results. Assisting with translation is not applicable to this work. Additionally, we used generative AI tools for creating or modifying scientific figures or images, editing to improve readability, identifying relevant literature. We have reviewed all AI-assisted work; for example, we checked all LLM-generated writings and codes, as well as the suggested related literature. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

Implementation details for optimization and inference are provided in Appendices A and B, with proofs in Appendix C. Appendix D describes the data and simulation settings. We will release our code together with a synthetic risk-score generator calibrated to the housing data, which are not publicly available. The experiment in Appendix E can be reproduced entirely from public data.

REFERENCES

Chunrong Ai and Xiaohong Chen. Efficient estimation of models with conditional moment restrictions containing unknown functions. *Econometrica*, 71(6):1795–1843, 2003.

Allegheny County Department of Human Services. Allegheny Housing Assessment (AHA): Frequently Asked Questions (FAQs), August 2020. URL https://analytics.alleghenycounty.us/wp-content/uploads/2020/08/20-ACDHS-16-ALPHA-FAQs-v6_2.pdf.

Nick Arnosti and Peng Shi. Design of lotteries and wait-lists for affordable housing allocation. *Management Science*, 66(6):2291–2307, 2020.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. Design-based identification with formula instruments: A review. *The Econometrics Journal*, 28(1):83–108, 2025.

Gary Chamberlain. Efficiency bounds for semiparametric regression. *Econometrica: Journal of the Econometric Society*, pp. 567–596, 1992.

Yash Chandak, Shiv Shankar, Vasilis Syrgkanis, and Emma Brunskill. Adaptive instrument design for indirect experiments. In *International Conference on Learning Representations*, volume 2024, pp. 46639–46688, 2024.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018. ISSN 1368-4221. doi: 10.1111/ectj.12097. URL https://doi.org/10.1111/ectj.12097.

M Colvin-Adams, M Valapour, M Hertz, B Heubner, K Paulson, V Dhungel, MA Skeans, L Edwards, V Ghimire, C Waller, et al. Lung and heart allocation in the united states. *American Journal of Transplantation*, 12(12):3213–3234, 2012.

Clément De Chaisemartin and Luc Behaghel. Estimating the effect of treatments allocated by randomized waiting lists. *Econometrica*, 88(4):1453–1477, 2020.

Joy Frechtling. The 2010 user-friendly handbook for project evaluation. Division of Research and Learning in Formal and Informal Settings, National Science Foundation, 2010. URL https://www.purdue.edu/research/oevprp/docs/pdf/2010NSFuser-friendlyhandbookforprojectevaluation.pdf.

Jinyong Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, pp. 315–331, 1998.

Jinyong Hahn, Keisuke Hirano, and Dean Karlan. Adaptive experimental design using the propensity score. *Journal of Business & Economic Statistics*, 29(1):96–108, 2011.

Guido W Imbens and Jeffrey M Wooldridge. Recent developments in the econometrics of program evaluation. *Journal of economic literature*, 47(1):5–86, 2009.

Masahiro Kato, Takuya Ishihara, Junya Honda, and Yusuke Narita. Efficient adaptive experimental design for average treatment effect estimation, 2025. URL https://arxiv.org/abs/2002.05308.

Edward H. Kennedy. Nonparametric Causal Effects Based on Incremental Propensity Score Interventions. *Journal of the American Statistical Association*, 114(526):645–656, April 2019. ISSN 0162-1459. doi: 10.1080/01621459.2017.1422737. URL https://doi.org/10.1080/01621459.2017.1422737. _eprint: https://doi.org/10.1080/01621459.2017.1422737.

Edward H Kennedy. Semiparametric doubly robust targeted double machine learning: a review. *Handbook of statistical methods for precision medicine*, pp. 207–236, 2024.

Jacob D Leshno. Dynamic matching in overloaded waiting lists. *American Economic Review*, 112 (12):3876–3910, 2022.

Medicaid and CHIP Payment and Access Commission (MACPAC). State management of home- and community-based services waiver waiting lists. Technical report, MACPAC, 8 2020. URL https://www.macpac.gov/wp-content/uploads/2020/08/State-Management-of-Home-and-Community-Based-Services-Waiver-Waiting-Lists.pdf.

Zeky Murra-Anton and Neil Thakral. The public housing allocation problem. *Mimeo*, 2026.

Jerzy Neyman. On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science*, 5(4):465–472, 1990. doi: 10.1214/ss/1177012031. Translated and edited by D. M. Dabrowska and T. P. Speed from the 1923 Polish original.

Miruna Oprescu, Brian Cho, and Nathan Kallus. Efficient adaptive experimentation with noncompliance. *Advances in Neural Information Processing Systems*, 38:42462–42498, 2026.

Peter M Robinson. Root-n-consistent semiparametric regression. *Econometrica: journal of the Econometric Society*, pp. 931–954, 1988.

Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688, 1974.

James Schummer. Influencing waiting lists. *Journal of Economic Theory*, 195:105263, 2021.

Hao Sun, Evan Munro, Georgy Kalashnov, Shuyang Du, and Stefan Wager. Treatment allocation under uncertain costs. *arXiv preprint arXiv:2103.11066*, 2021.

Max Tabord-Meehan. Stratification trees for adaptive randomisation in randomised controlled trials. *Review of Economic Studies*, 90(5):2646–2673, 2023.

Bill Tang, Çağıl Koçyiğit, Eric Rice, and Phebe Vayanos. Learning optimal and fair policies for online allocation of scarce societal resources from data collected in deployment. *arXiv preprint arXiv:2311.13765*, 2023.

Neil Thakral. Matching with stochastic arrival. In *AEA Papers and Proceedings*, volume 109, pp. 209–212. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2019.

Aad W. Van Der Vaart and Jon A. Wellner. *Weak Convergence and Empirical Processes*. Springer Series in Statistics. Springer New York, New York, NY, 1996. ISBN 978-1-4757-2547-6 978-1-4757-2545-2. doi: 10.1007/978-1-4757-2545-2. URL http://link.springer.com/10.1007/978-1-4757-2545-2.

Virginia Department of Education. Paying for child care. Virginia Child Care Website, 2024. URL https://www.childcare.virginia.gov/families/paying-for-child-care.

Virginia Department of Education. Stitch by stich: Building strong ccsp case management essentials! BPRO Fall 2025 Training Presentation, 2025. URL https://www.mybpro.org/uploads/1/4/0/9/140955706/child_care_subsidy_program-_building_strong_ccsp_case_management_essentials_1.pdf.

Bryan Wilder and Pim Welle. Learning treatment effects while treating those in need. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pp. 448–473, 2025.

Khurram Yamin, Edward Kennedy, and Bryan Wilder. Dependent randomized rounding for budget constrained experimental design. *arXiv preprint arXiv:2506.12677*, 2025.

Angela Zhou. Mind the gap: Optimal and equitable encouragement policies. *Available at SSRN 6630658*, 2023.

A Optimization details

Recall that we let $\eta \in \mathcal{U}$ denote the set of unknown nuisance functions at the time of optimization. Then, our optimization problem is

$$\min_{\theta:\mathcal{X}\to\Delta_k} \max_{\eta\in\mathcal{U}} \mathbb{V}_\psi(\theta,\eta) \quad \text{s.t. } \mathbb{E}[\theta(X)] = p;\ \mathbb{E}[\pi(X;\theta)u(X)] \geq c. \tag{3}$$

It is worth noting that the above optimization problem can be easily extended to include other constraints, such as bounding the maximum difference in treatment probabilities between demographic groups to satisfy fairness objectives. Formally, following Wilder & Welle (2025), if each additional constraint j is of the form $\mathbb{E}\{g_j(\pi(X;\theta),X)\} \leq c_j$, as long as the constraint is convex in θ it may be added to the optimization problem in (3).

In the sections that follow, we provide the complete sample optimization objective under both the exogenous and arrival settings. This objective is generally not strictly convex in θ, and so we conclude by discussing the choice of strictly convex regularization terms $R(\theta)$ which may be added to the objective to obtain unique solutions.

A.1 Optimization under exogenous arrivals

Recall the asymptotic variance from Theorem 5:

$$\lim_{n\to\infty} n\text{Var}(\widehat{\psi}^{\text{DR}}_{\text{ATE}}) = \mathbb{E}\left\{\frac{\text{Var}(Y \mid Z=1,X)}{\pi(X;\theta)} + \frac{\text{Var}(Y \mid Z=0,X)}{1-\pi(X;\theta)}\right\} + \text{Var}(\mu_1(X)-\mu_0(X)).$$

Under exogenous arrivals, we assume that $\eta = (\text{Var}(Y \mid Z=1,X), \text{Var}(Y \mid Z=0,X)) \equiv (\eta_1,\eta_0)$. It may be the case that the policymaker has estimates and/or reasonable assumptions about these nuisances, for example using historical data to estimate $\text{Var}(Y \mid Z=0,X)$ and assuming $\text{Var}(Y \mid Z=1,X)$ is bounded above by some function $h(\text{Var}(Y \mid Z=0,X),X)$. In this case, the policymaker can plug-in these values directly into the optimization problem. If the policymaker does not provide these values, we assume that the uncertainty set is of the form $\mathcal{U} = \{\eta : \|\eta\|_\infty \leq \delta\}$, i.e. that all nuisances are bounded above. In this case, the objective in (3) under the limiting asymptotic variance in Theorem 5 may be rewritten as:

$$\begin{aligned}
&\min_{\theta:\mathcal{X}\to\Delta_k} \max_{\eta\in\mathcal{U}} \mathbb{E}\left\{\frac{\eta_1}{\pi(X;\theta)} + \frac{\eta_0}{1-\pi(X;\theta)}\right\} + \text{Var}(\mu_1(X)-\mu_0(X)) \\
&= \min_{\theta:\mathcal{X}\to\Delta_k} \delta\mathbb{E}\left\{\frac{1}{\pi(X;\theta)} + \frac{1}{1-\pi(X;\theta)}\right\}.
\end{aligned}$$

And hence, the full optimization problem simplifies to

$$\begin{aligned}
&\min_{\theta:\mathcal{X}\to\Delta_k} \mathbb{E}\left\{\frac{1}{\pi(X;\theta)}+\frac{1}{1-\pi(X;\theta)}\right\}\\
&\pi(X;\theta)=\sum_{k=1}^{K}\alpha_k(\beta,p)\theta_k(X),\\
&\mathbb{E}\left\{\theta(X)\right\}=p,\\
&\mathbb{E}\left\{\pi(X)u(X)\right\}\geq c.
\end{aligned}$$

As in Wilder & Welle (2025), we reformulate the above optimization problem in terms of the dual to yield a separable inner objective:

$$\max_{\lambda\geq 0}\mathbb{E}\left\{\min_{\theta(X)\to\Delta_k}\frac{1}{\pi(X;\theta)}+\frac{1}{1-\pi(X;\theta)}+\sum_{j=1}^{J}\lambda_j(g(\theta(X),X)-c_j)\right\},$$

where λ_j is the dual variable associated with constraint j. The program now may be solved using the sample analogue

$$\max_{\lambda\geq 0}\frac{1}{n}\sum_{i=1}^{n}\left\{\min_{\theta(X_i)\to\Delta_k}\frac{1}{\pi(X_i;\theta)}+\frac{1}{1-\pi(X_i;\theta)}+\sum_{j=1}^{J}\lambda_j(g(\theta(X_i),X_i)-c_j)\right\}.$$

As all constraints are affine and the objective is convex in θ, the problem is convex and so solutions may be obtained using standard solvers.

A.2 Optimization under endogenous arrivals

Analogously, we recall the asymptotic variance under Theorem 6:

$$\lim_{n\to\infty} n\text{Var}(\widehat{\psi}_Z)=\left(\mathbb{E}\left[\frac{\{\alpha_Q(\beta,p)-\pi(X,\theta)\}^2}{\sigma(X)}\right]\right)^{-1}.$$

In the case of endogenous arrivals, $\eta=\sigma(X)=\mathbb{E}\left(U^2\mid X\right)$. We assume that the policymaker does not have access to this value beforehand, but as in the exogenous arrivals case that the uncertainty set is of the form $\mathcal{U}=\{\eta:\|\eta\|_\infty\leq\delta)$. And so, the objective in (3) simplifies to

$$\min_{\theta:\mathcal{X}\to\Delta_k}\max_{\eta\in\mathcal{U}}\left(\mathbb{E}\left[\frac{\{\alpha_Q(\beta,p)-\pi(X,\theta)\}^2}{\eta}\right]\right)^{-1}=\min_{\theta:\mathcal{X}\to\Delta_k}\left(\frac{1}{\delta}\mathbb{E}\left[\{\alpha_Q(\beta,p)-\pi(X,\theta)\}^2\right]\right)^{-1}.$$

Because δ is constant, the optimization problem under endogenous arrivals is thus

$$\begin{aligned}
&\min_{\theta:\mathcal{X}\to\Delta_k}\frac{1}{\mathbb{E}[\{\alpha_Q(\beta,p)-\pi(X,\theta)\}^2]}\\
&\pi(X;\theta)=\sum_{k=1}^{K}\alpha_k(\beta,p)\theta_k(X),\\
&\mathbb{E}\left[\theta(X)\right]=p,\\
&\mathbb{E}\left[\pi(X)u(X)\right]\geq c.
\end{aligned}$$

Unlike the exogenous arrivals case, the objective is not separable when formulated as a minimization problem. Notice that we may consider an equivalent objective of the form

$$\max_{\theta:\mathcal{X}\to\Delta_k}\mathbb{E}\left[\{\alpha_Q(\beta,p)-\pi(X;\theta)\}^2\right],$$

which may be rewritten as

$$\max_{\theta:\mathcal{X}\to\Delta_k}\mathbb{E}\left[\left\{\sum_{k=1}^{K}\alpha_k(\beta,p)^2\theta_k(X)-\left(\sum_{k=1}^{K}\alpha_k(\beta,p)\theta_k(X)\right)^2\right\}\right].$$

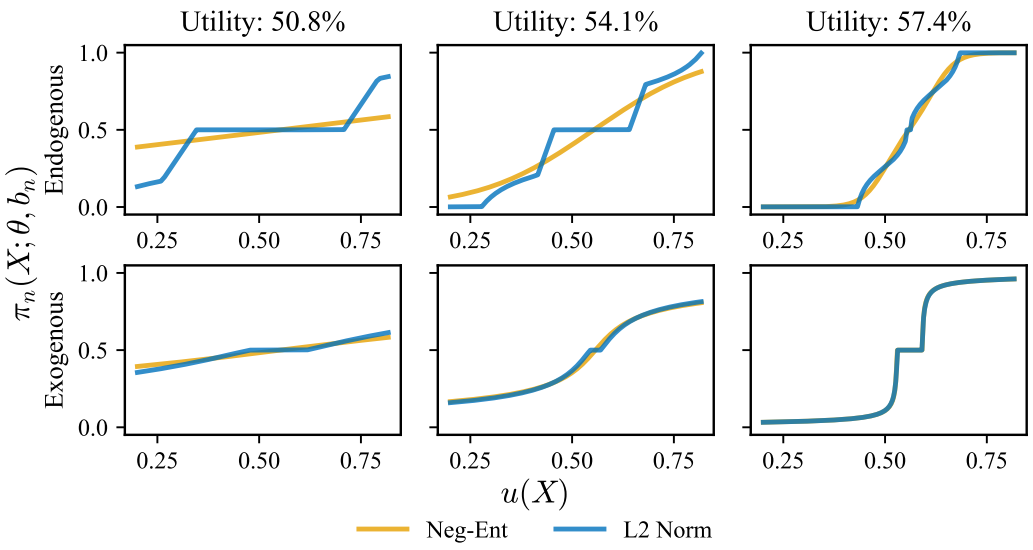

Figure 3: Effect of Negative-Entropy (orange) vs. L2 (blue) regularization on treatment probabilities $\pi_n(X;\theta,b_n)$ by individual utility $u(X)$ across varying utility constraints. L2-regularization consistently favors more deterministic solutions, assigning low treatment probability to low-utility individuals and high to high-utility, with the difference most pronounced in the endogenous setting at low levels of the utility constraint.

The objective is thus to maximize the expected conditional variance induced by the randomization policy θ, which is concave in θ. We may now rewrite the objective in terms of the dual

$$\min_{\lambda\geq 0}\max_{\theta:\mathcal{X}\to\Delta_k}\mathbb{E}\left[\left\{\sum_{k=1}^{K}\alpha_k(\beta,p)^2\theta_k(X)-\left(\sum_{k=1}^{K}\alpha_k(\beta,p)\theta_k(X)\right)^2\right\}-\sum_{j=1}^{J}\lambda_j(g(\theta(X),X)-c_j)\right],$$

where we once again use λ_j to denote the dual variable associated with constraint j. This is now separable, and so we can solve the problem using the sample population

$$\min_{\lambda\geq 0}\frac{1}{n}\sum_{i=1}^{n}\max_{\theta(X_i)\to\Delta_k}\left\{\sum_{k=1}^{K}\alpha_k(\beta,p)^2\theta_k(X_i)-\left(\sum_{k=1}^{K}\alpha_k(\beta,p)\theta_k(X_i)\right)^2\right\}-\sum_{j=1}^{J}\lambda_j(g(\theta(X_i),X_i)-c_j).$$

Solutions may be obtained using standard solvers as in the exogenous case.

A.3 Regularization

As π is affine in θ, neither objective is not strictly convex and thus there may be multiple valid solutions of θ. To obtain a unique solution, we augment the objective with a strictly convex regularizer $R(\theta)$ that selects a single optimizer among potentially many designs that obtain the same value of the primary efficiency objective. The choice of $R(\theta)$ can additionally encode a secondary intent; for example, a negative-entropy regularizer $R(\theta)=-\mathbb{E}\left[\sum_{k=1}^{K}\theta_k(X)\log(\theta_k(X))\right]$ favors higher-randomization queue assignments, while a quadratic regularizer such as $R(\theta)=\mathbb{E}\left[\|\theta(X)-p\|_2^2\right]$ promotes stability by discouraging deviations from the baseline proportions p. We demonstrate how the choice of regularization affects individual treatment probabilities in Figure 3.

B How should experiments be analyzed?

As alluded in Section 4, once the experiment is complete and the outcomes are observed, one can conduct standard statistical inference by constructing the Wald confidence interval $\widehat{\psi}\pm 1.96\sqrt{\widehat{\mathrm{Var}}(\widehat{\psi})}$

for $\psi \in \{\psi_{\text{ATE}}, \psi_Z\}$ based on the following asymptotic normality results. Here $\widehat{\text{Var}}(\widehat{\psi})$ denotes any consistent estimator of the asymptotic variance, such as the empirical variance of the estimated EIF divided by n.

Theorem 8 (Asymptotic normality of DR estimator)**.** *Assume the conditions of Theorem 5 and additionally that there exists $m' < \infty$ such that $\mathbb{E}(|Y_1|^4 + |Y^0|^4) < m'$. Then conditional on $\mathcal{D}$,*

$$\sqrt{n}\big(\widehat{\psi}_{\text{ATE}}^{\text{DR}} - \psi_{\text{ATE}}\big) \rightsquigarrow \mathcal{N}\left(0, V(\theta)\right),$$

where

$$V(\theta) \equiv \mathbb{E}\left\{ \frac{\text{Var}(Y^1 \mid X_i)}{\pi(X;\theta)} + \frac{\text{Var}(Y^0 \mid X_i)}{1-\pi(X;\theta)} \right\} + \text{Var}\big(\mu_1(X) - \mu_0(X)\big).$$

Corollary 1 (Asymptotic normality of $\widehat{\psi}_Z^\star$)**.** *Assume the conditions of Theorem 6 and additionally that there exists $m' < \infty$ such that $\mathbb{E}(U^4) < m'$. Then conditional on $\mathcal{D}$,*

$$\sqrt{n}\big(\widehat{\psi}_Z^\star - \psi_Z\big) \rightsquigarrow \mathcal{N}(0, V_Z(\theta)), \qquad V_Z(\theta) \equiv \left(\mathbb{E}\left[\frac{\{\alpha_Q(\beta,p) - \pi(X;\theta)\}^2}{\sigma(X)} \right] \right)^{-1}.$$

C Proofs

In this section, we provide proofs for the results discussed in both the main text and the appendix. Throughout the proofs, for a generic random variable V, we use the notation $V^n \equiv (V_1, \ldots, V_n)$ and adopt shorthands $\pi_n(x) = \pi_n(x;\theta,b_n)$ and $\tilde{\pi}_n(x,q) = \tilde{\pi}_n(x,q;\theta,b_n)$ for notational convenience. We also implicitly condition on the separate independent sample $\mathcal{D}$ in the proofs of Theorems 5, 6, 8, and Corollary 1.

While we suppressed unit indices throughout the main text since the data are iid, we introduce indices such as i and j in the proofs when needed to distinguish units or to analyze dependence induced by the queueing process.

C.1 Proof of Theorem 1

Proof. Fix $i \in \{1, \ldots, n\}$ and let $O_{-i} = \{(X_j, A_j, Q_j) : j \neq i\}$. Then $Z_i = \Phi(A_i, Q_i, O_{-i}, b_n)$ for some measurable function Φ that encodes our design.

By the iid assumption, together with independent queue assignment across units, O_{-i} is independent of $(X_i, A_i, Q_i, Y_i^0, Y_i^1)$. Therefore

$$O_{-i} \perp\!\!\!\perp (A_i, Q_i, Y_i^0, Y_i^1) \mid X_i.$$

Further, $Q_i \perp\!\!\!\perp (A_i, Y_i^0, Y_i^1) \mid X_i$ by the queue assignment mechanism and $A_i \perp\!\!\!\perp (Y_i^0, Y_i^1) \mid X_i$ by assumption. Hence

$$(A_i, Q_i) \perp\!\!\!\perp (Y_i^0, Y_i^1) \mid X_i.$$

It follows that

$$(A_i, Q_i, O_{-i}) \perp\!\!\!\perp (Y_i^0, Y_i^1) \mid X_i.$$

Hence for any bounded measurable h, we have

$$\begin{aligned}
\mathbb{E}\{h(Z_i) \mid X_i, Y_i^0, Y_i^1\} &= \mathbb{E}\{h(\Phi(A_i, Q_i, O_{-i}, b_n)) \mid X_i, Y_i^0, Y_i^1\} \\
&= \mathbb{E}\{h(\Phi(A_i, Q_i, O_{-i}, b_n)) \mid X_i\} \\
&= \mathbb{E}\{h(Z_i) \mid X_i\},
\end{aligned}$$

which is equivalent to

$$Z_i \perp\!\!\!\perp (Y_i^0, Y_i^1) \mid X_i.$$

Since the choice of i was arbitrary, this proves conditional randomization.

Given conditional randomization, identification is straightforward. Fix any $n \geq n_0$. As an example, consider the ATE:

$$\begin{aligned}
\mathbb{E}(Y_i^1 - Y_i^0) &= \mathbb{E}\{\mathbb{E}(Y_i^1 \mid X_i) - \mathbb{E}(Y_i^0 \mid X_i)\} \\
&= \mathbb{E}\{\mathbb{E}(Y_i^1 \mid Z_i = 1, X_i) - \mathbb{E}(Y_i^0 \mid Z_i = 0, X_i)\} \\
&= \mathbb{E}\{\mathbb{E}(Y_i \mid Z_i = 1, X_i) - \mathbb{E}(Y_i \mid Z_i = 0, X_i)\},
\end{aligned}$$

where the second equality follows from conditional randomization and the third from consistency. Positivity allows conditional means to be well-defined under the n-unit design. □

C.2 Proof of Lemma 1

Proof. Let Z^k denote the potential treatment status if we set $Q = k$. Further, let Y^{zk} denote the potential outcome if we set $Z = z$ and $Q = k$.

First consider conditional IV-unconfoundedness. Since queue assignments are conditionally randomized given X and $r_n(X, Q)$ is a function only of (X, Q), it follows that

$$r_n(X, Q) \perp\!\!\!\perp \{(Y^{0k}, Y^{1k}, Z^k) : k = 1, \ldots, K\} \mid X.$$

Next consider exclusion. By design, queue assignment only determines priority order for treatments and does not alter any other program feature. Therefore, under our stationary potential outcome formulation in which outcomes depend on whether treatment is eventually received by the horizon, but not on the precise timing of receipt, queue assignment affects outcomes only through terminal treatment status:

$$Y^{zk} = Y^z \quad \text{for all } z \in \{0, 1\},\ k \in \{1, \ldots, K\}.$$

Then any function of the queue assignment, such as $r_n(X, Q)$, can affect the observed outcome only through Z. Hence $r_n(X, Q)$ also satisfies exclusion.

Finally, note that monotonicity based on queue assignments

$$Z^k \geq Z^\ell \quad \text{for all } k < \ell,$$

also follows straightforwardly from our design; moving a unit to a higher-priority queue cannot reduce its treatment status. Hence taking conditional expectations yields

$$\tilde{\pi}_n(x, k) = \mathbb{E}(Z^k \mid X = x) \geq \mathbb{E}(Z^\ell \mid X = x) = \tilde{\pi}_n(x, \ell) \quad \text{for all } k < \ell. \qquad \square$$

C.3 Proof of Theorem 2

Proof. First, note that we can write

$$Y = Y^0 + Z^Q(Y^1 - Y^0).$$

Consider $\mathbb{E}\{r_n(X, Q)Y \mid X = x\}$. Substituting the above yields

$$\begin{aligned}
\mathbb{E}\{r_n(X, Q)Y \mid X = x\} &= \mathbb{E}\Big[(r_n(X, Q)\{Y^0 + Z^Q(Y^1 - Y^0)\} \mid X = x\Big] \\
&= \mathbb{E}\{r_n(X, Q)Y^0 \mid X = x\} + \mathbb{E}\Big[(r_n(X, Q)Z^Q(Y^1 - Y^0) \mid X = x\Big] \\
&= 0 + \sum_{k=1}^{K} \theta_k(x)\{\tilde{\pi}_n(x, k) - \pi_n(x)\}\mathbb{E}\{Z^k(Y^1 - Y^0) \mid X = x\} \\
&= \sum_{k<\ell} \theta_k(x)\theta_\ell(x)\{\tilde{\pi}_n(x, k) - \tilde{\pi}_n(x, \ell)\}\mathbb{E}\{(Z^k - Z^\ell)(Y^1 - Y^0) \mid X = x\},
\end{aligned}$$

where the second equality holds from $\mathbb{E}\{r_n(X, Q) \mid X\} = 0$ and conditional unconfoundeness.

Now recall that we have $Z^k \geq Z^\ell$ for $k < \ell$ by monotonicity, which implies $Z^k - Z^\ell = \mathbb{1}(i \in \mathcal{C}_{k,\ell})$. It follows that

$$\begin{aligned}
\mathbb{E}\{(Z^k - Z^\ell)(Y^1 - Y^0) \mid X = x\} &= \mathbb{P}(i \in \mathcal{C}_{k,\ell} \mid X = x)\mathbb{E}(Y^1 - Y^0 \mid X = x, i \in \mathcal{C}_{k,\ell}) \\
&= \{\tilde{\pi}_n(x, k) - \tilde{\pi}_n(x, \ell)\}\psi_{k,\ell}(x).
\end{aligned}$$

Hence

$$\mathbb{E}\{r_n(X, Q)Y \mid X = x\} = \sum_{k<\ell} \theta_k(x)\theta_\ell(x)\{\tilde{\pi}_n(x, k) - \tilde{\pi}_n(x, \ell)\}^2\psi_{k,\ell}(x),$$

and subsequently

$$\mathbb{E}\{r_n(X, Q)Y\} = \mathbb{E}\left\{\sum_{k<\ell} w_{k,\ell}(X)\psi_{k,\ell}(X)\right\}.$$

Now consider $\mathbb{E}\{r_n(X,Q)Z \mid X=x\}$. Since $Z = Z^Q$, an analogous derivation yields

$$\begin{aligned}\mathbb{E}\{r_n(X,Q)Z \mid X=x\} &= \sum_{k=1}^{K} \theta_k(x)\{\tilde{\pi}_n(x,k) - \pi_n(x)\}\mathbb{E}\{Z^k \mid X=x\} \\ &= \sum_{k<\ell} \theta_k(x)\theta_\ell(x)\{\tilde{\pi}_n(x,k) - \tilde{\pi}_n(x,\ell)\}^2.\end{aligned}$$

Hence $\mathbb{E}\{r_n(X,Q)Z\} = \mathbb{E}\{\sum_{k<\ell} w_{k,\ell}(X)\}$, and we have the desired. □

C.4 Proof of Theorem 3

Proof. This is a direct consequence of Theorem 6. □

C.5 Proof of Theorem 4

To show Theorem 4 holds, we first prove a lemma that shows the number of the eventually treated units in queues $1,\dots,k$ converges to $\min\{\beta, c_k\}$ defined in the statement of Theorem 4 under Assumptions 1–3, which is an intuitive result under our design.

Lemma 2. *Suppose Assumptions 1–3 hold. Let $Z_{i,t}$ be the treatment status of unit i at time t where $Z_{i,\tau} = Z_i$. For each queue $k \in \{1,\dots,K\}$ and each treatment allocation time $t \in \{1,\dots,\tau\}$, define*

$$\bar{T}_{\le k,t,n} \equiv \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}(Q_i \le k, Z_{i,t} = 1).$$

Then, as $n \to \infty$,

$$\sup_{1\le t\le \tau} \left|\bar{T}_{\le k,t,n} - \min\{\beta, c_k\}\mathbb{P}(A_i \le t)\right| \to 0 \quad \textit{a.s.}$$

In particular,

$$\frac{T_{\le k,n}}{n} \to \min\{\beta, c_k\} \quad \textit{a.s.}, \quad T_{\le k,n} \equiv \sum_{i=1}^{n} \mathbb{1}(Q_i \le k, Z_i = 1).$$

Proof. For each $k \in \{1,\dots,K\}$ and $t \in \{1,\dots,\tau\}$, define

$$\begin{aligned}\bar{N}_{\le k,\le t,n} &\equiv \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}(Q_i \le k, A_i \le t), \\ \bar{W}_{\le k,\le t,n} &\equiv \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}(Q_i \le k, A_i \le t, Z_{i,t} = 0), \\ \bar{b}_{\le t,n} &\equiv \frac{1}{n}\sum_{s=1}^{t} b_{s,n},\end{aligned}$$

which denotes, on $1/n$ scale, the average number of units that arrive in queues up to queue k by period-t allocation, the average number of units that arrive in queues up to queue k by period-t allocation and have not been treated, and the average budget up to period-t allocation, respectively, with the conventions

$$\bar{N}_{\le k,\le 0,n} = 0, \quad \bar{b}_{\le 0,n} = 0, \quad \bar{W}_{\le k,\le 0,n} = 0.$$

Further, define

$$\Delta\bar{N}_{\le k,t,n} \equiv \bar{N}_{\le k,\le t,n} - \bar{N}_{\le k,\le t-1,n},$$

which corresponds to the units that arrive after period-$t-1$ allocation and before period-t allocation. Our first goal is to show

$$\bar{W}_{\le k,\le t,n} = \max_{0\le s\le t}\Big\{\big(\bar{N}_{\le k,\le t,n} - \bar{N}_{\le k,\le s,n}\big) - \big(\bar{b}_{\le t,n} - \bar{b}_{\le s,n}\big)\Big\}, \tag{4}$$

for every t. Fix k and consider the quantity

$$\bar{W}_{\le k,\le t-1,n} + \Delta\bar{N}_{\le k,t,n},$$

which corresponds to the average number of waiting units in queues up to queue k *immediately* before the period-t allocation. Since the average number of period-t treatment slots used on queues $1, \ldots, k$ is exactly

$$\min\left\{\bar{W}_{\leq k,\leq t-1,n} + \Delta\bar{N}_{\leq k,t,n}, b_{t,n}/n\right\},$$

by our design, we can write

$$\begin{aligned}\bar{W}_{\leq k,\leq t,n} &= \bar{W}_{\leq k,\leq t-1,n} + \Delta\bar{N}_{\leq k,t,n} - \min\{\bar{W}_{\leq k,\leq t-1,n} + \Delta\bar{N}_{\leq k,t,n}, b_{t,n}/n\} \\ &= \left(\bar{W}_{\leq k,\leq t-1,n} + \Delta\bar{N}_{\leq k,t,n} - b_{t,n}/n\right)_+.\end{aligned}$$

Next we use induction to prove equation 4. For the base case $t = 0$, both sides are zero. Now suppose equation 4 holds for $t-1$. Then

$$\begin{aligned}\bar{W}_{\leq k,\leq t,n} &= \left(\max_{0\leq s\leq t-1}\left\{\left(\bar{N}_{\leq k,\leq t-1,n} - \bar{N}_{\leq k,\leq s,n}\right) - \left(\bar{b}_{\leq t-1,n} - \bar{b}_{\leq s,n}\right)\right\} + \Delta\bar{N}_{\leq k,t,n} - b_{t,n}/n\right)_+ \\ &= \left(\max_{0\leq s\leq t-1}\left\{\left(\bar{N}_{\leq k,\leq t,n} - \bar{N}_{\leq k,\leq s,n}\right) - \left(\bar{b}_{\leq t,n} - \bar{b}_{\leq s,n}\right)\right\}\right)_+ \\ &= \max_{0\leq s\leq t}\left\{\left(\bar{N}_{\leq k,\leq t,n} - \bar{N}_{\leq k,\leq s,n}\right) - \left(\bar{b}_{\leq t,n} - \bar{b}_{\leq s,n}\right)\right\},\end{aligned}$$

where the second equality follows from $\bar{N}_{\leq k,\leq t-1,n}+\Delta\bar{N}_{\leq k,t,n} = \bar{N}_{\leq k,\leq t,n}$ and $\bar{b}_{\leq t-1,n}+b_{t,n}/n = \bar{b}_{\leq t,n}$, and the third from the fact that the term corresponding to $s = t$ equals 0.

Now let $a_t \equiv \mathbb{P}(A_i \leq t)$. As a second step, we show

$$\sup_{1\leq t\leq\tau}\left|\bar{W}_{\leq k,\leq t,n} - \max_{0\leq s\leq t}(c_k-\beta)(a_t-a_s)\right| \to 0,$$

and

$$\max_{0\leq s\leq t}(c_k-\beta)(a_t-a_s) = (c_k-\beta)_+ a_t,$$

to arrive at

$$\sup_{1\leq t\leq\tau}\left|\bar{W}_{\leq k,\leq t,n} - (c_k-\beta)_+ a_t\right| \to 0.$$

Recall that Assumption 2 states

$$\sup_{1\leq t\leq\tau}\left|\bar{b}_{\leq t,n} - \beta a_t\right| \to 0,$$

and Assumptions 1 and 3 along with strong law of large numbers give

$$\bar{N}_{\leq k,\leq t,n} = \frac{1}{n}\sum_{i=1}^{n}\mathbb{1}(Q_i \leq k, A_i \leq t) \to c_k a_t \quad \text{a.s.}$$

As a result,

$$\sup_{1\leq t\leq\tau}\left|\bar{N}_{\leq k,\leq t,n} - c_k a_t\right| + \sup_{1\leq t\leq\tau}\left|\bar{b}_{\leq t,n} - \beta a_t\right| \to 0 \quad \text{a.s.}$$

Now for every $0 \leq s \leq t \leq \tau$, note that

$$\begin{aligned}&\left|\left(\bar{N}_{\leq k,\leq t,n} - \bar{N}_{\leq k,\leq s,n}\right) - \left(\bar{b}_{\leq t,n} - \bar{b}_{\leq s,n}\right) - (c_k-\beta)(a_t-a_s)\right| \\ &\leq \left|\bar{N}_{\leq k,\leq t,n} - c_k a_t\right| + \left|\bar{N}_{\leq k,\leq s,n} - c_k a_s\right| + \left|\bar{b}_{\leq t,n} - \beta a_t\right| + \left|\bar{b}_{\leq s,n} - \beta a_s\right| \\ &\leq 2\left\{\sup_{1\leq t\leq\tau}\left|\bar{N}_{\leq k,\leq t,n} - c_k a_t\right| + \sup_{1\leq t\leq\tau}\left|\bar{b}_{\leq t,n} - \beta a_t\right|\right\}.\end{aligned}$$

Taking the maximum over $0 \leq s \leq t$ gives

$$\sup_{1\leq t\leq\tau}\left|\bar{W}_{\leq k,\leq t,n} - \max_{0\leq s\leq t}(c_k-\beta)(a_t-a_s)\right| \to 0.$$

Since a_t is non-decreasing in t and $a_0 = 0$, there are two cases:

1. If $c_k \geq \beta$, then $(c_k - \beta)(a_t - a_s)$ is maximized at $s = 0$, and
$$\max_{0 \leq s \leq t} (c_k - \beta)(a_t - a_s) = (c_k - \beta)a_t.$$

2. If $c_k < \beta$, then $(c_k - \beta)(a_t - a_s) \leq 0$ for all $s \leq t$, so
$$\max_{0 \leq s \leq t} (c_k - \beta)(a_t - a_s) = 0.$$

It follows that
$$\sup_{1 \leq t \leq \tau} \left| \bar{W}_{\leq k, \leq t, n} - (c_k - \beta)_+ a_t \right| \to 0.$$

Finally, since
$$\bar{T}_{\leq k,t,n} = \bar{N}_{\leq k, \leq t, n} - \bar{W}_{\leq k, \leq t, n},$$
by construction, we have
$$\begin{aligned}
\sup_{1 \leq t \leq \tau} \left| \bar{T}_{\leq k,t,n} - \min\{\beta, c_k\} a_t \right| &= \sup_{1 \leq t \leq \tau} \left| \bar{T}_{\leq k,t,n} - \left\{ c_k a_t - (c_k - \beta)_+ a_t \right\} \right| \\
&= \sup_{1 \leq t \leq \tau} \left| \bar{N}_{\leq k, \leq t, n} - \bar{W}_{\leq k, \leq t, n} - \left\{ c_k a_t - (c_k - \beta)_+ a_t \right\} \right| \\
&\leq \sup_{1 \leq t \leq \tau} \left| \bar{N}_{\leq k, \leq t, n} - c_k a_t \right| + \sup_{1 \leq t \leq \tau} \left| \bar{W}_{\leq k, \leq t, n} - (c_k - \beta)_+ a_t \right| \to 0.
\end{aligned}$$

Taking $t = \tau$ and using $a_\tau = 1$ gives $T_{\leq k,n}/n = \bar{T}_{\leq k,\tau,n} \to \min\{\beta, c_k\}$. □

Proof of Theorem 4. all statements involving conditioning on $X_i = x$ should be understood to hold for almost every x, and all convergences below are almost sure convergences except for the deterministic budget path.

It suffices to show that, for each k with $p_k > 0$ and each x such that $\theta_k(x) > 0$,
$$\tilde{\pi}_n(x, k) = \mathbb{P}_{\theta, b_n}(Z_i = 1 \mid X_i = x, Q_i = k) \to \alpha_k(\beta, p).$$
Fix such a k and consider
$$T_{k,n} \equiv \sum_{i=1}^{n} \mathbb{1}(Q_i = k, Z_i = 1), \quad N_{k,n} \equiv \sum_{i=1}^{n} \mathbb{1}(Q_i = k),$$
which represent the number of individuals in queue k who are treated by the end of the allocation process, and the total number of individuals that end up in queue k, respectively.

Let $J_k \equiv \{j : Q_j = k\}$ where $|J_k| = N_{k,n}$. Conditional on $Q^n = q^n$ with $q_i = k$, the set J_k is fixed. Since the arrival times are iid, continuous, and independent of (X_i, Q^n), the arrival order within J_k is a uniform random permutation even after conditioning on $X_i = x$. Under our treatment allocation rule, for any realization of the arrivals, exactly $T_{k,n}$ of the $N_{k,n}$ units in queue k are treated by the end of the allocation process. Therefore, we have
$$\tilde{\pi}_n(x, k) = \mathbb{E}\left(\frac{T_{k,n}}{N_{k,n}} \,\middle|\, X_i = x, Q_i = k \right).$$

Now consider
$$T_{\leq k,n} \equiv \sum_{i=1}^{n} \mathbb{1}(Q_i \leq k, Z_i = 1), \quad T_{\leq 0,n} \equiv 0,$$
as defined in Lemma 2 and note that $T_{k,n} = T_{\leq k,n} - T_{\leq k-1,n}$. It follows by Lemma 2 that
$$\begin{aligned}
\frac{T_{k,n}}{n} &= \frac{T_{\leq k,n} - T_{\leq k-1,n}}{n} \to \min\{\beta, c_k\} - \min\{\beta, c_{k-1}\} \\
&= \beta - (\beta - c_k)_+ - \{\beta - (\beta - c_{k-1})_+\} \\
&= (\beta - c_{k-1})_+ - (\beta - c_k)_+,
\end{aligned}$$
and
$$\frac{T_{k,n}}{N_{k,n}} = \frac{T_{k,n}/n}{N_{k,n}/n} \to \frac{(\beta - c_{k-1})_+ - (\beta - c_k)_+}{p_k} = \alpha_k(\beta, p).$$
Therefore, dominated convergence implies
$$\mathbb{E}\left(\frac{T_{k,n}}{N_{k,n}} \,\middle|\, X_i = x, Q_i = k \right) \to \alpha_k(\beta, p).$$ □

C.6 Proof of Theorem 5

We first present two lemmas that are necessary to show Theorem 5. The first lemma says that changing one unit's queue assignment or arrival time can change the number of eventually treated units in any queue-period cell by at most a constant. The second lemma is a more technical one that shows the quantity

$$\left|\mathbb{E}\Big[\{Z_i - \pi_n(X_i)\}\{Z_j - \pi_n(X_j)\} \mid X_i, X_j\Big]\right|,$$

is of order n^{-1}.

Throughout the proofs of the two lemmas, we write $S_i = s$ to denote $A_i \in [s-1, s)$ and $C_i = (Q_i, S_i) \in \mathcal{C} = \{1, \ldots, K\} \times \{1, \ldots, \tau\}$.

Lemma 3. *Define $V \equiv \{(Q_1, A_1), \ldots, (Q_n, A_n)\}$. For each $r \in \{1, \ldots, n\}$, let $V^{(r)}$ be the vector obtained from V by replacing its r-th coordinate (Q_r, A_r) by an arbitrary (Q_r', A_r'). For $c = (k, s) \in \mathcal{C}$, define*

$$T_{c,n}(V) \equiv \sum_{i=1}^{n} \mathbb{1}(C_i = c,\ Z_i(V) = 1).$$

Then there exists a constant $L = L(K, \tau) < \infty$ such that, for every $r \in \{1, \ldots, n\}$,

$$\sup_{c \in \mathcal{C}} \left|T_{c,n}(V) - T_{c,n}(V^{(r)})\right| \le L.$$

Proof. For $c = (k, s) \in \mathcal{C}$ and $t \in \{1, \ldots, \tau\}$, define

$$W_{c,\le t,n}(V) \equiv \sum_{i=1}^{n} \mathbb{1}(C_i = c,\ Z_{i,t-1}(V) = 0,\ s \le t),$$

$$D_{c,t,n}(V) \equiv \sum_{i=1}^{n} \mathbb{1}(C_i = c,\ Z_{i,t-1}(V) = 0,\ Z_{i,t}(V) = 1),$$

where $W_{c,\le t,n}(V)$ is the number of untreated units from cell $c = (k, s)$ available for treatment at the beginning of period t, and $D_{c,t,n}(V)$ is the number newly treated from cell c in period t. By construction,

$$T_{c,n}(V) = \sum_{t=s}^{\tau} D_{c,t,n}(V).$$

By our treatment rule, cells are served in lexicographic order: first by queue k, then by arrival period s. Hence for $s > t$, $D_{c,t,n}(V) = 0$. For $c = (k, s)$ with $s \le t$, define

$$R^{-}_{c,t,n}(V) \equiv b_{t,n} - \sum_{\ell<k} \sum_{u \le t} W_{(\ell,u),\le t,n}(V) - \sum_{u<s} W_{(k,u),\le t,n}(V),$$

and

$$R^{+}_{c,t,n}(V) \equiv b_{t,n} - \sum_{\ell<k} \sum_{u \le t} W_{(\ell,u),\le t,n}(V) - \sum_{u\le s} W_{(k,u),\le t,n}(V),$$

as the the remaining budget after serving all cells that have higher priority than (k, s), and the remaining budget after *also* serving cell (k, s), respectively. Then

$$D_{c,t,n}(V) = \{R^{-}_{c,t,n}(V)\}_+ - \{R^{+}_{c,t,n}(V)\}_+.$$

Now let

$$\Delta_t \equiv \sum_{c \in \mathcal{C}} \left|W_{c,\le t,n}(V) - W_{c,\le t,n}(V^{(r)})\right|,$$

which measures the total discrepancy in the number of units who have arrived by period t and have not yet been treated under V versus $V^{(r)}$. We use Δ_t to upper bound $|D_{c,t,n}(V) - D_{c,t,n}(V^{(r)})|$, and show that Δ_t is bounded by some finite constant $C(K, \tau)$.

Suppose $t = 1$. Changing one coordinate of V can remove at most one unit from one cell and add at most one unit to another, so $\Delta_1 \le 2$.

Next, fix $t \in \{1, \dots, \tau - 1\}$. If $s > t$, then $D_{c,t,n}(V) = D_{c,t,n}(V^{(r)}) = 0$. If $s \le t$, then

$$\begin{aligned}
&\left|D_{c,t,n}(V) - D_{c,t,n}(V^{(r)})\right| \\
&= \left|\{R^-_{c,t,n}(V)\}_+ - \{R^+_{c,t,n}(V)\}_+ - \{R^-_{c,t,n}(V^{(r)})\}_+ + \{R^+_{c,t,n}(V^{(r)})\}_+\right| \\
&\le \left|\{R^-_{c,t,n}(V)\}_+ - \{R^-_{c,t,n}(V^{(r)})\}_+\right| + \left|\{R^+_{c,t,n}(V)\}_+ - \{R^+_{c,t,n}(V^{(r)})\}_+\right| \\
&\le \left|R^-_{c,t,n}(V) - R^-_{c,t,n}(V^{(r)})\right| + \left|R^+_{c,t,n}(V) - R^+_{c,t,n}(V^{(r)})\right|.
\end{aligned}$$

where the last line follows from $u \mapsto u_+$ being 1-Lipschitz. Further note that

$$\left|R^-_{c,t,n}(V) - R^-_{c,t,n}(V^{(r)})\right| \le \sum_{c' \in \mathcal{C}} \left|W_{c',\le t,n}(V) - W_{c',\le t,n}(V^{(r)})\right| = \Delta_t,$$

and similarly $\left|R^+_{c,t,n}(V) - R^+_{c,t,n}(V^{(r)})\right| \le \Delta_t$. Therefore, for every $c \in \mathcal{C}$,

$$\left|D_{c,t,n}(V) - D_{c,t,n}(V^{(r)})\right| \le 2\Delta_t.$$

Summing over $c \in \mathcal{C}$ gives

$$\sum_{c \in \mathcal{C}} \left|D_{c,t,n}(V) - D_{c,t,n}(V^{(r)})\right| \le \sum_{c \in \mathcal{C}} 2\Delta_t = 2K\tau\,\Delta_t.$$

Now let $A^{\text{new}}_{c,t+1,n}(V)$ denote the number of units from cell c that become newly available in period $t+1$. Then we can write

$$W_{c,\le t+1,n}(V) = W_{c,\le t,n}(V) - D_{c,t,n}(V) + A^{\text{new}}_{c,t+1,n}(V).$$

Now triangle inequality implies

$$\begin{aligned}
\Delta_{t+1} &= \sum_{c \in \mathcal{C}} \left|W_{c,\le t+1,n}(V) - W_{c,\le t+1,n}(V^{(r)})\right| \\
&\le \sum_{c \in \mathcal{C}} \left|W_{c,\le t,n}(V) - W_{c,\le t,n}(V^{(r)})\right| + \sum_{c \in \mathcal{C}} \left|D_{c,t,n}(V) - D_{c,t,n}(V^{(r)})\right| + \sum_{c \in \mathcal{C}} \left|A^{\text{new}}_{c,t+1,n}(V) - A^{\text{new}}_{c,t+1,n}(V^{(r)})\right| \\
&\le (1 + 2K\tau)\Delta_t + \sum_{c \in \mathcal{C}} \left|A^{\text{new}}_{c,t+1,n}(V) - A^{\text{new}}_{c,t+1,n}(v^{(r)})\right|.
\end{aligned}$$

Since V and $V^{(r)}$ differ in only one coordinate, the newly available arrival-count vectors in period $t+1$ differ in ℓ_1 distance by at most 2. Thus

$$\Delta_{t+1} \le \Delta_t + 2K\tau\,\Delta_t + 2 = (1 + 2K\tau)\Delta_t + 2,$$

and subsequently, after some algebra, $\Delta_t \le 2\sum_{h=0}^{t-1}(1 + 2K\tau)^h$. Since τ is fixed, we obtain

$$\max_{1 \le t \le \tau} \Delta_t \le 2\sum_{h=0}^{\tau-1}(1 + 2K\tau)^h \equiv C(K, \tau)$$

Consequently, $\left|D_{c,t,n}(V) - D_{c,t,n}(V^{(r)})\right| \le 2C(K,\tau)$ for every $c \in \mathcal{C}$ and every t.

Finally, for $c = (k, s)$,

$$\begin{aligned}
\left|T_{c,n}(V) - T_{c,n}(V^{(r)})\right| &= \left|\sum_{t=s}^{\tau}\{D_{c,t,n}(V) - D_{c,t,n}(V^{(r)})\}\right| \\
&\le \sum_{t=s}^{\tau}\left|D_{c,t,n}(V) - D_{c,t,n}(V^{(r)})\right| \\
&\le 2\tau C(K,\tau).
\end{aligned}$$

Taking the supremum over $c \in \mathcal{C}$, the result holds with $L(K,\tau) = 2\tau C(K,\tau)$. □

Lemma 4. *Suppose Assumption 3 holds. For $c = (k, s) \in \mathcal{C}$, let $\lambda_{\min} \equiv \min\{\lambda_c : \lambda_c > 0\} > 0$ where $\lambda_c \equiv \mathbb{P}(C_i = c) = p_k \mathbb{P}(S_i = s)$. Then there exists $C = C(\lambda_{\min}, L) < \infty$ such that, for all $n \geq 2$ and all $i \neq j$,*

$$\left| \mathbb{E}\Big[\{Z_i - \pi_n(X_i)\}\{Z_j - \pi_n(X_j)\} \mid X_i, X_j \Big] \right| \leq \frac{C}{n} \quad \text{a.s.}$$

Proof. Note that

$$\begin{aligned} &\mathbb{E}\big[\{Z_i - \pi_n(X_i)\}\{Z_j - \pi_n(X_j)\} \mid X_i, X_j\big] \\ &= \text{Cov}(Z_i, Z_j \mid X_i, X_j) \qquad (5) \\ &\quad + \{\mathbb{E}(Z_i \mid X_i, X_j) - \pi_n(X_i)\}\{\mathbb{E}(Z_j \mid X_i, X_j) - \pi_n(X_j)\}. \qquad (6) \end{aligned}$$

The goal of our proof is to show that equation 5 is $O(n^{-1})$ and equation 6 is $O(n^{-2})$.

First we show equation 5 is $O(n^{-1})$. For each $c \in \mathcal{C}$, define

$$N_{c,n} \equiv \sum_{i=1}^{n} \mathbb{1}(C_i = c), \qquad T_{c,n} \equiv \sum_{i=1}^{n} \mathbb{1}(C_i = c, Z_i = 1),$$

which represent the total and the treated number of units in cell c. Further, define $m_c(C^n) \equiv T_{c,n}/N_{c,n}$ whenever $N_{c,n} > 0$. Conditional on X_i, X_j, C^n, as in the proof of Theorem 4, we have

$$\mathbb{E}(Z_i \mid X_i, X_j, C^n) = m_{C_i}(C^n), \qquad \mathbb{E}(Z_j \mid X_i, X_j, C^n) = m_{C_j}(C^n).$$

So by the law of total covariance,

$$\begin{aligned} &\text{Cov}(Z_i, Z_j \mid X_i, X_j) \\ &= \mathbb{E}\{\text{Cov}(Z_i, Z_j \mid X_i, X_j, C^n) \mid X_i, X_j\} + \text{Cov}\{\mathbb{E}(Z_i \mid X_i, X_j, C^n), \mathbb{E}(Z_j \mid X_i, X_j, C^n) \mid X_i, X_j\} \\ &= \mathbb{E}\{\text{Cov}(Z_i, Z_j \mid X_i, X_j, C^n) \mid X_i, X_j\} + \text{Cov}\{m_{C_i}(C^n), m_{C_j}(C^n) \mid X_i, X_j\}. \qquad (7) \end{aligned}$$

For the first term in equation 7, note that $\text{Cov}(Z_i, Z_j \mid X_i, X_j, C^n) = 0$ when $C_i \neq C_j$ since, conditional on C^n, the remaining randomness comes from independent within-cell arrival orders. In the case $C_i = C_j = c$, sampling without replacement gives

$$\begin{aligned} |\text{Cov}(Z_i, Z_j \mid X_i, X_j, C^n)| &= |\mathbb{E}(Z_i Z_j \mid X_i, X_j, C^n) - \mathbb{E}(Z_i \mid X_i, X_j, C^n)\mathbb{E}(Z_j \mid X_i, X_j, C^n)| \\ &= \frac{(T_{c,n}/N_{c,n})(1 - T_{c,n}/N_{c,n})}{N_{c,n} - 1} \leq \frac{1}{4(N_{c,n} - 1)}. \end{aligned}$$

It follows that

$$|\mathbb{E}\{\text{Cov}(Z_i, Z_j \mid X_i, X_j, C^n) \mid X_i, X_j\}| \leq \frac{1}{4}\mathbb{E}\left\{ \frac{\mathbb{1}(C_i = C_j)}{N_{C_i,n} - 1} \;\middle|\; X_i, X_j \right\}.$$

Now fix $c \in \mathcal{C}$ with $\lambda_c > 0$. Conditional on X_i, X_j and $C_i = C_j = c$, we claim that the random variables $\mathbb{1}(C_m = c)$, $m \neq i, j$, remain iid Bernoulli with probability λ_c. To see this, note that the units are iid across m, and the queue assignments and arrival times are generated independently across units. Therefore, for every $m \neq i, j$, C_m is independent of (X_i, X_j, C_i, C_j). Hence conditioning on X_i, X_j and on the event $C_i = C_j = c$ does not change the distribution of C_m, and

$$\mathbb{P}(C_m = c \mid X_i, X_j, C_i = C_j = c) = \mathbb{P}(C_m = c) = \lambda_c.$$

Moreover, the variables $\{C_m : m \neq i, j\}$ remain mutually independent under this conditioning. Consequently,

$$N_{c,n} = 2 + \sum_{m \neq i,j} \mathbb{1}(C_m = c) \equiv 2 + B, \qquad B \sim \text{Bin}(n - 2, \lambda_c).$$

It follows that

$$\begin{aligned} \mathbb{E}\left(\frac{1}{N_{c,n} - 1} \;\middle|\; C_i = C_j = c \right) &= \mathbb{E}\left(\frac{1}{B + 1} \right) \\ &= \sum_{b=0}^{n-2} \binom{n-2}{b} \lambda_c^b (1 - \lambda_c)^{n-2-b} \int_0^1 t^b dt \\ &= \frac{1 - (1 - \lambda_c)^{n-1}}{(n-1)\lambda_c} \leq \frac{1}{(n-1)\lambda_c} \leq \frac{1}{(n-1)\lambda_{\min}}. \end{aligned}$$

Therefore

$$|\mathbb{E}\{\mathrm{Cov}(Z_i, Z_j \mid X_i, X_j, C^n) \mid X_i, X_j\}| \le \frac{1}{4(n-1)\lambda_{\min}}.$$

Now we consider the second term in equation 7. Write

$$\begin{aligned}
&\mathrm{Cov}\{m_{C_i}(C^n), m_{C_j}(C^n) \mid X_i, X_j\} \\
&= \mathbb{E}\left[\mathrm{Cov}\{m_{C_i}(C^n), m_{C_j}(C^n) \mid X_i, X_j, C_i, C_j\} \mid X_i, X_j\right] \qquad (8)\\
&+ \mathrm{Cov}\left(\mathbb{E}\{m_{C_i}(C^n) \mid X_i, X_j, C_i, C_j\}, \mathbb{E}\{m_{C_j}(C^n) \mid X_i, X_j, C_i, C_j\} \mid X_i, X_j\right). \qquad (9)
\end{aligned}$$

We first show that equation 8 is $O(n^{-1})$ by using Cauchy–Schwarz and applying Efron–Stein inequality. Consider the event $C_i = c$ and $C_j = d$, and define $V \equiv (C_m)_{m \neq i,j}$.Then $m_c(C^n)$ can be viewed as a function of V with $C_i = c$ and $C_j = d$ fixed, and we write it as $m_c(V)$. For each $r \in \{1, \dots, n-2\}$, let $V^{(r)}$ be the vector obtained from V by replacing its r-th coordinate by an independent copy C'_r with the same distribution as C_r. Changing one coordinate of V changes one unit's queue-period cell, which can be realized by changing that unit's (Q, A) coordinate; hence Lemma 3 applies.

Then for each $c' \in \mathcal{C}$, we have

$$|N_{c',n}(V) - N_{c',n}(V^{(r)})| \le 1.$$

Moreover, by Lemma 3,

$$|T_{c',n}(V) - T_{c',n}(V^{(r)})| \le L.$$

It follows that

$$\begin{aligned}
|m_c(V) - m_c(V^{(r)})| &= \left|\frac{T_{c,n}(V)}{N_{c,n}(V)} - \frac{T_{c,n}(V^{(r)})}{N_{c,n}(V^{(r)})}\right| \\
&= \frac{\left|T_{c,n}(V)N_{c,n}(V^{(r)}) - T_{c,n}(V^{(r)})N_{c,n}(V)\right|}{N_{c,n}(V)N_{c,n}(V^{(r)})} \\
&\le \frac{|T_{c,n}(V) - T_{c,n}(V^{(r)})|N_{c,n}(V^{(r)}) + T_{c,n}(V^{(r)})|N_{c,n}(V) - N_{c,n}(V^{(r)})|}{N_{c,n}(V)N_{c,n}(V^{(r)})} \\
&\le \frac{LN_{c,n}(V^{(r)}) + N_{c,n}(V^{(r)})}{N_{c,n}(V)N_{c,n}(V^{(r)})} = \frac{L+1}{N_{c,n}(V)}.
\end{aligned}$$

Now, note that $N_{c,n}(V) \ge 1 + B$ conditional on $C_i = c$ and $C_j = d$ where $B \sim \mathrm{Bin}(n-2, \lambda_c)$ as before. Therefore

$$\begin{aligned}
\mathbb{E}\left[\{m_c(V) - m_c(V^{(r)})\}^2 \mid X_i, X_j, C_i = c, C_j = d\right] &\le (L+1)^2\mathbb{E}\left(\frac{1}{(B+1)^2}\right) \\
&\le 2(L+1)^2\mathbb{E}\left(\frac{1}{(B+1)(B+2)}\right).
\end{aligned}$$

Now using $1/\{(b+1)(b+2)\} = \int_0^1 t^b(1-t)dt$ and the binomial theorem, we have

$$\begin{aligned}
\mathbb{E}\left\{\frac{1}{(B+1)(B+2)}\right\} &= \int_0^1 (1 - \lambda_c + \lambda_c t)^{n-2}(1-t)dt \\
&= \frac{1}{\lambda_c^2}\left\{\frac{1-(1-\lambda_c)^{n-1}}{n-1} - \frac{1-(1-\lambda_c)^n}{n}\right\} \le \frac{1}{n(n-1)\lambda_{\min}^2}.
\end{aligned}$$

Combining the above yields

$$\mathbb{E}\left[\{m_c(V) - m_c(V^{(r)})\}^2 \mid X_i, X_j, C_i = c, C_j = d\right] \le \frac{2(L+1)^2}{n(n-1)\lambda_{\min}^2}.$$

Therefore by Efron–Stein inequality,

$$\begin{aligned}
\mathrm{Var}(m_c(C^n) \mid X_i, X_j, C_i = c, C_j = d) &\le \frac{1}{2}\sum_{r=1}^{n-2}\mathbb{E}\left[\{m_c(V) - m_c(V^{(r)})\}^2 \mid X_i, X_j, C_i = c, C_j = d\right] \\
&\le \frac{n-2}{2} \cdot \frac{2(L+1)^2}{n(n-1)\lambda_{\min}^2} \le \frac{(L+1)^2}{n\lambda_{\min}^2}.
\end{aligned}$$

An identical bound holds for $\mathrm{Var}(m_d(C^n) \mid X_i, X_j, C_i = c, C_j = d)$. Hence by Cauchy–Schwarz,

$$|\mathrm{Cov}(m_c(C^n), m_d(C^n) \mid X_i, X_j, C_i = c, C_j = d)| \le \frac{(L+1)^2}{n\lambda_{\min}^2}.$$

uniformly over c and d, and subsequently,

$$\left|\mathbb{E}\left[\mathrm{Cov}\{m_{C_i}(C^n), m_{C_j}(C^n) \mid X_i, X_j, C_i, C_j\} \mid X_i, X_j\right]\right| \le \frac{(L+1)^2}{n\lambda_{\min}^2}.$$

This bounds equation 8.

To conclude bounding equation 5, we now bound equation 9. Write

$$a_{c,d} \equiv \mathbb{E}\{m_c(C^n) \mid X_i, X_j, C_i = c, C_j = d\}, \qquad b_{c,d} \equiv \mathbb{E}\{m_d(C^n) \mid X_i, X_j, C_i = c, C_j = d\}.$$

Then the term of interest can be written as $\mathrm{Cov}(a_{C_i,C_j}, b_{C_i,C_j} \mid X_i, X_j)$.

Fix c and $d, d' \in \mathcal{C}$. For a given realization $V = v$, let

$$T_{c,n}(v,d) \equiv T_{c,n}(C^n) \,\Big|\, X_i, X_j, C_i = c, C_j = d, V = v,$$

$$N_{c,n}(v,d) \equiv N_{c,n}(C^n) \,\Big|\, X_i, X_j, C_i = c, C_j = d, V = v,$$

and $m_c(v,d) \equiv T_{c,n}(v,d)/N_{c,n}(v,d)$ following our notation from before.

Recall that, if we change only the value of C_j while holding $V = v$ and $C_i = c$ fixed, then

$$|T_{c,n}(v,d) - T_{c,n}(v,d')| \le L, \qquad |N_{c,n}(v,d) - N_{c,n}(v,d')| \le 1.$$

Hence

$$|m_c(V,d) - m_c(V,d')| \le \frac{L+1}{\min\{N_{c,n}(V,d), N_{c,n}(V,d')\}}.$$

Also,

$$\min\{N_{c,n}(V,d), N_{c,n}(V,d')\} \ge 1 + \sum_{m \ne i,j} \mathbb{1}(C_m = c) = 1 + B,$$

where $B \sim \mathrm{Bin}(n-2, \lambda_c)$. It follows from the calculations above that

$$\sup_{d,d'} |a_{c,d} - a_{c,d'}| \le (L+1)\mathbb{E}\left(\frac{1}{B+1}\right) \le \frac{L+1}{(n-1)\lambda_{\min}}.$$

Similarly,

$$\sup_{c,c'} |b_{c,d} - b_{c',d}| \le \frac{L+1}{(n-1)\lambda_{\min}}.$$

Now let

$$\bar{a}_{C_i}(X_j) \equiv \sum_{d \in \mathcal{C}} P(C_j = d \mid X_j) a_{C_i,d}, \qquad \bar{b}_{C_j}(X_i) \equiv \sum_{c \in \mathcal{C}} P(C_i = c \mid X_i) b_{c,C_j}.$$

and write the term of our interest as

$$\begin{aligned}
&\mathrm{Cov}(a_{C_i,C_j}, b_{C_i,C_j} \mid X_i, X_j) \\
&= \mathrm{Cov}\left(a_{C_i,C_j} - \bar{a}_{C_i}(X_j) + \bar{a}_{C_i}(X_j), b_{C_i,C_j} - \bar{b}_{C_j}(X_i) + \bar{b}_{C_j}(X_i) \mid X_i, X_j\right) \\
&= \mathrm{Cov}\left(a_{C_i,C_j} - \bar{a}_{C_i}(X_j), b_{C_i,C_j} - \bar{b}_{C_j}(X_i) \mid X_i, X_j\right) \\
&+ \mathrm{Cov}\left(a_{C_i,C_j} - \bar{a}_{C_i}(X_j), \bar{b}_{C_j}(X_i) \mid X_i, X_j\right) + \mathrm{Cov}\left(\bar{a}_{C_i}(X_j), b_{C_i,C_j} - \bar{b}_{C_j}(X_i) \mid X_i, X_j\right),
\end{aligned}$$

where $\mathrm{Cov}\{\bar{a}_{C_i}(X_j), \bar{b}_{C_j}(X_i) \mid X_i, X_j\} = 0$ since C_i and C_j are independent given X_i, X_j. Also,

$$\mathbb{E}\{a_{C_i,C_j} - \bar{a}_{C_i}(X_j) \mid X_i, X_j\} = 0, \qquad \mathbb{E}\{b_{C_i,C_j} - \bar{b}_{C_j}(X_i) \mid X_i, X_j\} = 0,$$

and $0 \le \bar{a}_{C_i}(X_j), \bar{b}_{C_j}(X_i) \le 1$. Therefore,

$$\begin{aligned}
\left|\mathrm{Cov}(a_{C_i,C_j}, b_{C_i,C_j} \mid X_i, X_j)\right| \le\ & \mathbb{E}\{|a_{C_i,C_j} - \bar{a}_{C_i}(X_j)| \mid X_i, X_j\} \\
&+ \mathbb{E}\{|b_{C_i,C_j} - \bar{b}_{C_j}(X_i)| \mid X_i, X_j\} \\
&+ \mathbb{E}\{|a_{C_i,C_j} - \bar{a}_{C_i}(X_j)||b_{C_i,C_j} - \bar{b}_{C_j}(X_i)| \mid X_i, X_j\}.
\end{aligned}$$

Now

$$|a_{C_i,C_j} - \bar{a}_{C_i}(X_j)| \le \sup_{d,d'} |a_{C_i,d} - a_{C_i,d'}|,$$

which implies

$$\mathbb{E}\{|a_{C_i,C_j} - \bar{a}_{C_i}(X_j)| \mid X_i, X_j\} \le \frac{L+1}{(n-1)\lambda_{\min}}.$$

Similarly,

$$\mathbb{E}\{|b_{C_i,C_j} - \bar{b}_{C_j}(X_i)| \mid X_i, X_j\} \le \frac{L+1}{(n-1)\lambda_{\min}},$$

and

$$\mathbb{E}\{|a_{C_i,C_j} - \bar{a}_{C_i}(X_j)||b_{C_i,C_j} - \bar{b}_{C_j}(X_i)| \mid X_i, X_j\} \le \frac{(L+1)^2}{(n-1)^2\lambda_{\min}^2}.$$

Combining the above bounds yields

$$\left|\mathrm{Cov}(a_{C_i,C_j}, b_{C_i,C_j} \mid X_i, X_j)\right| \le \frac{2(L+1)}{(n-1)\lambda_{\min}} + \frac{(L+1)^2}{(n-1)^2\lambda_{\min}^2}.$$

This bounds equation 9. Hence the combined bound on equation 5 is $O(n^{-1})$.

As a final piece, we bound equation 6 and show it is $O(n^{-2})$. Using the same $a_{c,d}$ as above, iterated expectation gives

$$\mathbb{E}(Z_i \mid X_i, X_j) = \sum_{c\in\mathcal{C}} \sum_{d\in\mathcal{C}} P(C_i = c \mid X_i) P(C_j = d \mid X_j) a_{c,d},$$

and

$$\pi_n(X_i) = \sum_{c\in\mathcal{C}} \sum_{d\in\mathcal{C}} P(C_i = c \mid X_i) \lambda_d a_{c,d}.$$

It follows that

$$\mathbb{E}(Z_i \mid X_i, X_j) - \pi_n(X_i) = \sum_{c\in\mathcal{C}} P(C_i = c \mid X_i) \sum_{d\in\mathcal{C}} \{P(C_j = d \mid X_j) - \lambda_d\}(a_{c,d} - \tilde{a}_c),$$

where $\tilde{a}_c \equiv \sum_{d\in\mathcal{C}} \lambda_d a_{c,d}$. Hence,

$$\begin{aligned}
|\mathbb{E}(Z_i \mid X_i, X_j) - \pi_n(X_i)| &\le \sum_{c\in\mathcal{C}} P(C_i = c \mid X_i) \left(\sum_{d\in\mathcal{C}} |P(C_j = d \mid X_j) - \lambda_d| \right) \sup_{d,d'} |a_{c,d} - a_{c,d'}| \\
&\le 2 \sum_{c\in\mathcal{C}} P(C_i = c \mid X_i) \sup_{d,d'} |a_{c,d} - a_{c,d'}| \le \frac{2(L+1)}{(n-1)\lambda_{\min}},
\end{aligned}$$

where the second inequality follows because $P(C_j = \cdot \mid X_j)$ and $\lambda_\cdot$ are probability vectors. Therefore, we also have

$$|\mathbb{E}(Z_j \mid X_i, X_j) - \pi_n(X_j)| \le \frac{2(L+1)}{(n-1)\lambda_{\min}}.$$

Consequently,

$$|\{\mathbb{E}(Z_i \mid X_i, X_j) - \pi_n(X_i)\}\{\mathbb{E}(Z_j \mid X_i, X_j) - \pi_n(X_j)\}| \le \frac{4(L+1)^2}{(n-1)^2\lambda_{\min}^2}.$$

This shows that equation 6 is $O(n^{-2})$, which completes the proof. □

Proof of Theorem 5. Throughout the proof, we assume $n \ge n_0$ for $n_0 \in \mathbb{N}$ so that positivity in Theorem 1 holds. Define the *oracle* influence function as

$$\phi_i \equiv \mu_1(X_i) - \mu_0(X_i) + \frac{Z_i}{\pi_n(X_i)}\{Y_i - \mu_1(X_i)\} - \frac{1 - Z_i}{1 - \pi_n(X_i)}\{Y_i - \mu_0(X_i)\}.$$

We first decompose $\widehat{\psi}_{\text{ATE}}^{\text{DR}}$ with respect to ϕ_i. Let

$$g(x) \equiv \frac{\widehat{\mu}_1(x) - \mu_1(x)}{\pi_n(x;\theta)} + \frac{\widehat{\mu}_0(x) - \mu_0(x)}{1 - \pi_n(x;\theta)}.$$

Some algebra gives

$$\widehat{\psi}_{\text{ATE}}^{\text{DR}} = \frac{1}{n}\sum_{i=1}^{n}\phi_i - \frac{1}{n}\sum_{i=1}^{n}\{Z_i - \pi_n(X_i)\}g(X_i) \equiv \frac{1}{n}\sum_{i=1}^{n}\phi_i - R_n,$$

so we have

$$n\text{Var}(\widehat{\psi}_{\text{ATE}}^{\text{DR}}) = n\text{Var}\left(\frac{1}{n}\sum_{i=1}^{n}\phi_i\right) - 2n\text{Cov}\left(\frac{1}{n}\sum_{i=1}^{n}\phi_i, R_n\right) + n\text{Var}(R_n). \tag{10}$$

We first show that $n\text{Var}(R_n) = o_{\mathbb{P}}(1)$. Note that

$$\mathbb{E}(R_n) = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\big[g(X_i)\{\mathbb{E}(Z_i \mid X_i) - \pi_n(X_i)\}\big] = 0,$$

so $\text{Var}(R_n) = \mathbb{E}(R_n^2)$. Now write

$$n\mathbb{E}(R_n^2) = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\big[g(X_i)^2\{Z_i - \pi_n(X_i)\}^2\big] + \frac{2}{n}\sum_{i<j}\mathbb{E}\big[g(X_i)g(X_j)\{Z_i - \pi_n(X_i)\}\{Z_j - \pi_n(X_j)\}\big].$$

For the diagonal term, $\{Z_i - \pi_n(X_i)\}^2 \le 1$ implies

$$\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\big[g(X_i)^2\{Z_i - \pi_n(X_i)\}^2\big] \le \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\{g(X_i)^2\} = \mathbb{E}\{g(X)^2\}.$$

For the off-diagonal term, using Lemma 4 and $2|ab| \le a^2 + b^2$, can be bounded as

$$\begin{aligned}
&\frac{2}{n}\sum_{i<j}\Big|\mathbb{E}\big[g(X_i)g(X_j)\{Z_i - \pi_n(X_i)\}\{Z_j - \pi_n(X_j)\}\big]\Big| \\
&\le \frac{2}{n}\sum_{i<j}\mathbb{E}\big(|g(X_i)g(X_j)|\big|\mathbb{E}\big[\{Z_i - \pi_n(X_i)\}\{Z_j - \pi_n(X_j)\} \mid X_i, X_j\big]\big|\big) \\
&\le \frac{2}{n}\sum_{i<j}\mathbb{E}\Big(|g(X_i)g(X_j)| \cdot \frac{C}{n}\Big) \le \frac{C}{n^2}\mathbb{E}\Big[\sum_{i<j}\{g(X_i)^2 + g(X_j)^2\}\Big] \le C\mathbb{E}\{g(X)^2\},
\end{aligned}$$

and therefore

$$n\mathbb{E}(R_n^2) \le (1 + C)\mathbb{E}\{g(X)^2\}.$$

Finally, positivity implies

$$g(X)^2 \le \frac{2}{\gamma^2}\Big[\{\widehat{\mu}_1(X) - \mu_1(X)\}^2 + \{\widehat{\mu}_0(X) - \mu_0(X)\}^2\Big],$$

so $\mathbb{E}\{g(X)^2\} = o_{\mathbb{P}}(1)$ by $\|\widehat{\mu}_z - \mu_z\|_{L^2(\mathbb{P})} = o_{\mathbb{P}}(1)$ for $z \in \{0, 1\}$. Thus we have $n\mathbb{E}(R_n^2) = o_{\mathbb{P}}(1)$.

Next we analyze the leading term in equation 10. Write

$$n\text{Var}\left(\frac{1}{n}\sum_{i=1}^{n}\phi_i\right) = \frac{1}{n}\sum_{i=1}^{n}\text{Var}(\phi_i) + \frac{2}{n}\sum_{i<j}\text{Cov}(\phi_i, \phi_j).$$

The analysis of the diagonal term is standard; by the law of total variance,

$$\text{Var}(\phi_i) = \mathbb{E}\{\text{Var}(\phi_i \mid X_i)\} + \text{Var}\{\mathbb{E}(\phi_i \mid X_i)\}.$$

Note that $\text{Var}\{\mathbb{E}(\phi_i \mid X_i)\} = \text{Var}(\mu_1(X_i) - \mu_0(X_i))$ since $\mathbb{E}(\phi_i \mid X_i) = \mu_1(X_i) - \mu_0(X_i)$, and

$$\text{Var}(\phi_i \mid X_i) = \frac{\text{Var}(Y_i \mid Z_i = 1, X_i)}{\pi_n(X_i)} + \frac{\text{Var}(Y_i \mid Z_i = 0, X_i)}{1 - \pi_n(X_i)}.$$

Therefore

$$\text{Var}(\phi_i) = \mathbb{E}\left\{\frac{\text{Var}(Y_i \mid Z_i = 1, X_i)}{\pi_n(X_i)} + \frac{\text{Var}(Y_i \mid Z_i = 0, X_i)}{1 - \pi_n(X_i)}\right\} + \text{Var}(\mu_1(X_i) - \mu_0(X_i)),$$

and applying Theorem 4 and dominated convergence under positivity yield

$$\mathbb{E}\left\{\frac{\text{Var}(Y_i \mid Z_i = 1, X_i)}{\pi_n(X_i)} + \frac{\text{Var}(Y_i \mid Z_i = 0, X_i)}{1 - \pi_n(X_i)}\right\} \to \mathbb{E}\left\{\frac{\text{Var}(Y_i \mid Z_i = 1, X_i)}{\pi(X_i;\theta)} + \frac{\text{Var}(Y_i \mid Z_i = 0, X_i)}{1 - \pi(X_i;\theta)}\right\}.$$

We now show that the covariance term is zero. Define

$$B_i \equiv \frac{Y_i^1 - \mu_1(X_i)}{\pi_n(X_i)} + \frac{Y_i^0 - \mu_0(X_i)}{1 - \pi_n(X_i)}.$$

Some algebra gives

$$\phi_i = (Y_i^1 - Y_i^0) + B_i(Z_i - \pi_n(X_i)).$$

Now fix $i \neq j$. Expanding the covariance, we have

$$\begin{aligned}
\text{Cov}(\phi_i, \phi_j) = \text{Cov}(Y_i^1 - Y_i^0, Y_j^1 - Y_j^0). &\quad (11)\\
+ \text{Cov}(Y_i^1 - Y_i^0, B_j\{Z_j - \pi_n(X_j)\}) &\quad (12)\\
+ \text{Cov}(B_i\{Z_i - \pi_n(X_i)\}, Y_j^1 - Y_j^0) &\quad (13)\\
+ \text{Cov}(B_i\{Z_i - \pi_n(X_i)\}, B_j\{Z_j - \pi_n(X_j)\}). &\quad (14)
\end{aligned}$$

equation 11 is zero by the iid assumption. We show that the remaining terms are also zero by conditioning on $H_n \equiv (X^n, A^n, Q^n)$.

Consider equation 12. We have

$$\begin{aligned}
\text{Cov}(Y_i^1 - Y_i^0, B_j\{Z_j - \pi_n(X_j)\}) = \mathbb{E}\left[\text{Cov}\big(Y_i^1 - Y_i^0, B_j\{Z_j - \pi_n(X_j)\} \mid H_n\big)\right]\\
+ \text{Cov}\left(\mathbb{E}(Y_i^1 - Y_i^0 \mid H_n), \mathbb{E}\big(B_j\{Z_j - \pi_n(X_j)\} \mid H_n\big)\right).
\end{aligned}$$

Note that $\{Z_j - \pi_n(X_j)\}$ is constant given H_n since Z^n is then a deterministic function of the design; further, $(Y_i^1 - Y_i^0) \perp\!\!\!\perp B_j$ since B_j is a function of Y_j^1, Y_j^0, and X_j. It follows that

$$\text{Cov}\big(Y_i^1 - Y_i^0, B_j\{Z_j - \pi_n(X_j)\} \mid H_n\big) = \{Z_j - \pi_n(X_j)\}\text{Cov}\big(Y_i^1 - Y_i^0, B_j \mid H_n\big) = 0.$$

Moreover, the covariance of the conditional means is also zero as

$$\mathbb{E}\big(B_j\{Z_j - \pi_n(X_j)\} \mid H_n\big) = \{Z_j - \pi_n(X_j)\}\mathbb{E}(B_j \mid H_n) = \{Z_j - \pi_n(X_j)\}\mathbb{E}(B_j \mid X_j) = 0.$$

Hence by the same logic, equation 13 is also zero.

Next, consider equation 14. Again, write

$$\begin{aligned}
&\text{Cov}(B_i\{Z_i - \pi_n(X_i)\}, B_j\{Z_j - \pi_n(X_j)\})\\
&\quad = \mathbb{E}\left[\text{Cov}\big(B_i\{Z_i - \pi_n(X_i)\}, B_j\{Z_j - \pi_n(X_j)\} \mid H_n\big)\right]\\
&\quad + \text{Cov}\left(\mathbb{E}\big(B_i\{Z_i - \pi_n(X_i)\} \mid H_n\big), \mathbb{E}\big(B_j\{Z_j - \pi_n(X_j)\} \mid H_n\big)\right).
\end{aligned}$$

$\{Z_i - \pi_n(X_i)\}$ and $\{Z_j - \pi_n(X_j)\}$ are fixed given H_n, so the conditional covariance reduces to

$$\{Z_i - \pi_n(X_i)\}\{Z_j - \pi_n(X_j)\}\text{Cov}(B_i, B_j \mid H_n) = 0,$$

since $B_i \perp\!\!\!\perp B_j$. Also,

$$\mathbb{E}\big(B_i\{Z_i - \pi_n(X_i)\} \mid H_n\big) = \{Z_i - \pi_n(X_i)\}\mathbb{E}(B_i \mid H_n) = 0,$$

and similarly for j, so the covariance of the conditional means is zero.

Combining the above, we conclude $\text{Cov}(\phi_i, \phi_j) = 0$ for all $i \neq j$. Therefore,

$$n\text{Var}\left(\frac{1}{n}\sum_{i=1}^n \phi_i\right) = \frac{1}{n}\sum_{i=1}^n \text{Var}(\phi_i),$$

The proof concludes by noting that the second term in equation 10 is

$$\left|2n\text{Cov}(n^{-1}\textstyle\sum_i \phi_i, R_n)\right| \leq 2\sqrt{n\text{Var}(n^{-1}\sum_i \phi_i)}\sqrt{n\text{Var}(R_n)} = o_{\mathbb{P}}(1),$$

by Cauchy-Schwarz, $n\text{Var}(n^{-1}\sum_i \phi_i) = O_{\mathbb{P}}(1)$, and $n\text{Var}(R_n) = o_{\mathbb{P}}(1)$. □

C.7 Proof of Theorem 6

We first shed light on how $\widehat{\psi}_Z^\star$ is motivated. Note that residualizing the PLIV model equation 1 with respect to X gives the conditional moment restriction

$$\mathbb{E}\Big[Y - m(X) - \psi_Z\{Z - \pi_n(X;\theta,b_n)\} \;\Big|\; X, Q\Big] = 0.$$

For any integrable function $f(X,Q)$, multiplying the above equation by $f(X,Q)$ yields the unconditional moment

$$\mathbb{E}\Big(f(X,Q)\big[Y - m(X) - \psi_Z\{Z - \pi_n(X;\theta,b_n)\}\big]\Big) = 0,$$

which implies

$$\psi_Z = \frac{\mathbb{E}\big[f(X,Q)\{Y - m(X)\}\big]}{\mathbb{E}\big[f(X,Q)\{Z - \pi_n(X;\theta,b_n)\}\big]},$$

provided the denominator is nonzero. Although any $f(X,Q)$ yields such an unconditional moment, only its residualized component with respect to X contributes after residualizing Y and Z. We therefore restrict without loss of generality to functions satisfying $\mathbb{E}\{f(X,Q) \mid X\} = 0$. For such functions,

$$\mathbb{E}\big[f(X,Q)\{Z - \pi_n(X;\theta,b_n)\}\big] = \mathbb{E}\big\{f(X,Q)Z\big\}.$$

Hence any $f(X,Q)$ satisfying $\mathbb{E}\{f(X,Q) \mid X\} = 0$ and $\mathbb{E}\{f(X,Q)Z\} \neq 0$ is a valid instrument under our queue-randomizing design. The corresponding sample analogue gives an estimator

$$\widehat{\psi}_Z(f) = \frac{\mathbb{P}_n\big[f(X,Q)\{Y - \widehat{m}(X)\}\big]}{\mathbb{P}_n\big[f(X,Q)\{Z - \pi_n(X;\theta,b_n)\}\big]}.$$

This formulation allows us to use the efficiency theory for semiparametric conditional moment restrictions (Chamberlain, 1992; Ai & Chen, 2003); although many instruments identify ψ_Z, only the optimal instrument attains the efficiency bound. In the theorem below, we first derive the asymptotic variance of $\widehat{\psi}_Z(f)$ for a generic uniformly bounded f satisfying $\mathbb{E}\{f(X,Q) \mid X\} = 0$ and $\mathbb{E}\{f(X,Q)Z\} \neq 0$, and then show that the oracle choice $f(X,Q) = r_n(X,Q)/\sigma(X)$ yields the iid efficiency bound, with the feasible estimator $\widehat{\psi}_Z^\star$ obtained by replacing $\sigma(X)$ with $\widehat{\sigma}(X)$ and attaining the same asymptotic variance.

Proof of Theorem 6. Throughout the proof, we write

$$f_i \equiv f(X_i,Q_i), \quad r_i \equiv r_n(X_i,Q_i), \quad \zeta_i \equiv \alpha_{Q_i}(\beta,p) - \pi(X_i;\theta), \quad \tilde{m}(x) \equiv \widehat{m}(x) - m(x),$$

for notational ease.

First, note that $m(X_i) = \psi_Z \pi_n(X_i) + g(X_i)$ since $\mathbb{E}(U_i \mid X_i) = 0$ by iterated expectation. Then we can write

$$\begin{aligned} Y_i - \widehat{m}(X_i) &= Y_i - m(X_i) + m(X_i) - \widehat{m}(X_i) \\ &= \psi_Z Z_i + g(X_i) + U_i - m(X_i) + m(X_i) - \widehat{m}(X_i) \\ &= \psi_Z\{Z_i - \pi_n(X_i)\} + U_i - \tilde{m}(X_i), \end{aligned}$$

which implies

$$\begin{aligned} \frac{1}{n}\sum_{i=1}^n f_i\{Y_i - \widehat{m}(X_i)\} &= \psi_Z \frac{1}{n}\sum_{i=1}^n f_i\{Z_i - \pi_n(X_i)\} + \frac{1}{n}\sum_{i=1}^n f_i U_i - \frac{1}{n}\sum_{i=1}^n f_i \tilde{m}(X_i) \\ &\equiv \psi_Z \Delta_n + S_n - R_n. \end{aligned}$$

Subsequently,

$$\widehat{\psi}_Z - \psi_Z = \frac{S_n - R_n}{\Delta_n}.$$

We first show $\Delta_n = \mathbb{E}(f_i\zeta_i) + o_{\mathbb{P}}(1) \equiv \Delta + o_{\mathbb{P}}(1)$. Decompose Δ_n as

$$\begin{aligned} \Delta_n &= \frac{1}{n}\sum_{i=1}^n f_i\{Z_i - \tilde{\pi}_n(X_i,Q_i) + \tilde{\pi}_n(X_i,Q_i) - \pi_n(X_i)\} \\ &= \frac{1}{n}\sum_{i=1}^n f_i\{Z_i - \tilde{\pi}_n(X_i,Q_i) + \frac{1}{n}\sum_{i=1}^n f_i r_i \equiv B_n + C_n. \end{aligned}$$

Consider B_n. Note that $\mathbb{E}(B_n) = 0$ since $\mathbb{E}\{Z_i - \tilde{\pi}_n(X_i, Q_i) \mid X_i, Q_i\} = 0$, so $n\text{Var}(B_n) = n\mathbb{E}(B_n^2)$. Now write

$$n\mathbb{E}(B_n^2) = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\Big[f_i^2\{Z_i - \tilde{\pi}_n(X_i, Q_i)\}^2\Big] + \frac{2}{n}\sum_{i<j}\mathbb{E}\Big[f_i f_j\{Z_i - \tilde{\pi}_n(X_i, Q_i)\}\{Z_j - \tilde{\pi}_n(X_j, Q_j)\}\Big].$$

The diagonal term is bounded by $\mathbb{E}(f_i^2)$ since $\{Z_i - \tilde{\pi}_n(X_i, Q_i)\}^2 \leq 1$. For the off-diagonal term, we introduce the following lemma, whose proof is similar to that of Lemma 4 and thus omitted.

Lemma 5. *There exists a constant $C > 0$ depending only on p such that for all $n \geq 2$ and all $i \neq j$,*

$$\Big|\mathbb{E}\Big[\{Z_i - \tilde{\pi}_n(X_i, Q_i)\}\{Z_j - \tilde{\pi}_n(X_j, Q_j)\}\Big|X_i, Q_i, X_j, Q_j\Big]\Big| \leq \frac{C}{n} \quad \textit{a.s.}$$

Lemma 5 and the uniform boundedness of f imply

$$\begin{aligned}
&\frac{2}{n}\sum_{i<j}\Big|\mathbb{E}\Big[f_i f_j\{Z_i - \tilde{\pi}_n(X_i, Q_i)\}\{Z_j - \tilde{\pi}_n(X_j, Q_j)\}\Big]\Big| \\
&\leq \frac{2}{n}\sum_{i<j}\mathbb{E}\Big(|f_i f_j|\Big|\mathbb{E}\Big[\{Z_i - \tilde{\pi}_n(X_i, Q_i)\}\{Z_j - \tilde{\pi}_n(X_j, Q_j)\}\Big|X_i, Q_i, X_j, Q_j\Big]\Big|\Big) \\
&\leq \frac{2c}{n^2}\sum_{i<j}\mathbb{E}|f_i f_j| = O(1).
\end{aligned}$$

Hence $n\text{Var}(B_n) = O(1)$ and, in turn, $B_n = O_{\mathbb{P}}(n^{-1/2}) = o_{\mathbb{P}}(1)$ by Chebyshev's inequality.

For C_n, write

$$C_n = \{C_n - \mathbb{E}(f_i r_i)\} + \mathbb{E}\{f_i(r_i - \zeta_i)\} + \Delta.$$

Since (X_i, Q_i) are iid across i, $f_i r_i$ are iid. Further, the fact that $|r_i| \leq 1$ and f_i is uniformly bounded implies $\text{Var}(f_i r_i)$ is also uniformly bounded. It thus follows that the first term is $o_{\mathbb{P}}(1)$. For the second term, $\mathbb{E}\{f_i(r_i - \zeta_i)\} \to 0$ a.s. by Theorem 4. It follows that

$$C_n = \Delta + o_{\mathbb{P}}(1),$$

and consequently $\Delta_n = B_n + C_n = \Delta + o_{\mathbb{P}}(1)$.

We next show

$$n\text{Var}(S_n - R_n) = n\text{Var}(S_n) + n\text{Var}(R_n) - 2n\text{Cov}(S_n, R_n) = \mathbb{E}(f_i^2 U_i^2) + o_{\mathbb{P}}(1).$$

First consider $n\text{Var}(S_n)$. Since $\mathbb{E}(f_i U_i) = 0$, we have

$$\begin{aligned}
n\text{Var}(S_n) &= \frac{1}{n}\sum_{i=1}^{n}\text{Var}(f_i U_i) + \frac{2}{n}\sum_{i<j}\text{Cov}(f_i U_i, f_j U_j) \\
&= \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}(f_i^2 U_i^2) + \frac{2}{n}\sum_{i<j}\text{Cov}(f_i U_i, f_j U_j) \\
&= \mathbb{E}(f_i^2 U_i^2),
\end{aligned}$$

where the third equality follows from

$$\text{Cov}(f_i U_i, f_j U_j) = \mathbb{E}(f_i U_i f_j U_j) = \mathbb{E}\{f_i f_j \mathbb{E}(U_i \mid X_i, Q_i)\mathbb{E}(U_j \mid X_j, Q_j)\} = 0.$$

Since $\tilde{m}$ is fixed conditional on $\mathcal{D}$, we can apply the same logic to show $n\text{Var}(R_n) = \text{Var}(f_i\tilde{m}(X_i))$. Then the uniform boundedness of f_i and $\|\widehat{m} - m\|_{L_2(\mathbb{P})} = o_{\mathbb{P}}(1)$ imply

$$n\text{Var}(R_n) = \text{Var}(f_i\tilde{m}(X_i)) \leq \mathbb{E}\{f_i^2\tilde{m}(X_i)^2\} \leq c'^2\|\tilde{m}\|_{L_2(\mathbb{P})}^2 = o_{\mathbb{P}}(1).$$

Further,

$$\text{Cov}(S_n, R_n) = \frac{1}{n^2}\sum_{i=1}^{n}\text{Cov}(f_i U_i, f_i\tilde{m}(X_i)) + \frac{2}{n^2}\sum_{i<j}\text{Cov}(f_i U_i, f_j\tilde{m}(X_j)) = 0,$$

since $\mathrm{Cov}(f_iU_i, f_i\tilde{m}(X_i)) = \mathbb{E}\{f_i^2\tilde{m}(X_i)\mathbb{E}(U_i \mid X_i, Q_i)\} = 0$ and $\mathrm{Cov}(f_iU_i, f_j\tilde{m}(X_j)) = 0$ by independence across units.

Putting what we have together, since $\Delta_n = \Delta + o_{\mathbb{P}}(1)$ and $n\mathrm{Var}(S_n - R_n) = O_{\mathbb{P}}(1)$, continuous mapping theorem and Chebyshev's inequality imply $\Delta_n^{-1} = \Delta^{-1} + o_{\mathbb{P}}(1)$ and $S_n - R_n = O_{\mathbb{P}}(n^{-1/2})$, respectively. It follows that

$$\frac{S_n - R_n}{\Delta_n} = \frac{S_n - R_n}{\Delta} + (S_n - R_n)\left(\frac{1}{\Delta_n} - \frac{1}{\Delta}\right) = \frac{S_n - R_n}{\Delta} + o_{\mathbb{P}}(n^{-1/2}),$$

and therefore

$$n\mathrm{Var}(\widehat{\psi}_Z) = n\mathrm{Var}(\widehat{\psi}_Z - \psi) = \frac{n\mathrm{Var}(S_n - R_n)}{\Delta^2} + o_{\mathbb{P}}(1) = \frac{\mathbb{E}(f_i^2U_i^2)}{\mathbb{E}(f_i\zeta_i)^2} + o_{\mathbb{P}}(1) = \frac{\mathbb{E}\{f_i^2\sigma(X_i)\}}{\mathbb{E}(f_i\zeta_i)^2} + o_{\mathbb{P}}(1),$$

where the last equality follows because, under our queue-randomizing design, Q_i is assigned using only X_i and independent design randomness. Hence $Q_i \perp\!\!\!\perp U_i \mid X_i$, so

$$\mathbb{E}(U_i^2 \mid X_i, Q_i) = \mathbb{E}(U_i^2 \mid X_i) = \sigma(X_i).$$

Now consider the oracle choice $f_i = r_i/\sigma(X_i)$. Applying Theorem 4 gives

$$\lim_{n\to\infty} n\mathrm{Var}\{\widehat{\psi}_Z)\} = \left(\mathbb{E}\left[\frac{\{\alpha_Q(\beta, p) - \pi(X;\theta)\}^2}{\sigma(X)}\right]\right)^{-1}.$$

Finally, replacing $\sigma(X)$ by $\widehat{\sigma}(X)$ does not change the limit under the assumed consistency of $\widehat{\sigma}$. Hence

$$\lim_{n\to\infty} n\mathrm{Var}(\widehat{\psi}_Z^\star) = \left(\mathbb{E}\left[\frac{\{\alpha_Q(\beta, p) - \pi(X;\theta)\}^2}{\sigma(X)}\right]\right)^{-1}. \qquad \square$$

C.8 Proof of Theorem 7

Proof. Refer to Section A for details. $\square$

C.9 Proof of Theorem 8

Proof. Throughout the proof, we assume $n \geq n_0$ for $n_0 \in \mathbb{N}$ to ensure positivity holds. As in the proof of Theorem 5, denote the *oracle* influence function as

$$\begin{aligned}\phi_i &\equiv \mu_1(X_i) - \mu_0(X_i) + \frac{Z_i}{\pi_n(X_i)}\{Y_i - \mu_1(X_i)\} - \frac{1 - Z_i}{1 - \pi_n(X_i)}\{Y_i - \mu_0(X_i)\} \\ &= \mu_1(X_i) - \mu_0(X_i) + \frac{Z_i}{\pi_n(X_i)}\{Y_i^1 - \mu_1(X_i)\} - \frac{1 - Z_i}{1 - \pi_n(X_i)}\{Y_i^0 - \mu_0(X_i)\},\end{aligned}$$

where the equality holds from consistency.

Since

$$\sqrt{n}\big(\widehat{\psi}_{\mathrm{ATE}}^{\mathrm{DR}} - \psi_{\mathrm{ATE}}\big) = \frac{1}{\sqrt{n}}\sum_{i=1}^n \big(\phi_i - \psi_{\mathrm{ATE}}\big) + o_{\mathbb{P}}(1),$$

it suffices to show that $n^{-1/2}\sum_{i=1}^n(\phi_i - \psi_{\mathrm{ATE}})$ converges in distribution to $\mathcal{N}(0, V(\theta))$.

Write

$$\begin{aligned}\phi_i - \psi_{\mathrm{ATE}} &= \big\{\mu_1(X_i) - \mu_0(X_i) - \psi_{\mathrm{ATE}}\big\} + \frac{Z_i}{\pi_n(X_i)}\{Y_i^1 - \mu_1(X_i)\} - \frac{1 - Z_i}{1 - \pi_n(X_i)}\{Y_i^0 - \mu_0(X_i)\} \\ &\equiv \big\{\mu_1(X_i) - \mu_0(X_i) - \psi_{\mathrm{ATE}}\big\} + R_{i,n}.\end{aligned}$$

Then

$$\frac{1}{\sqrt{n}}\sum_{i=1}^n(\phi_i - \psi_{\mathrm{ATE}}) = \frac{1}{\sqrt{n}}\sum_{i=1}^n R_{i,n} + \frac{1}{\sqrt{n}}\sum_{i=1}^n \big\{\mu_1(X_i) - \mu_0(X_i) - \psi_{\mathrm{ATE}}\big\} \equiv T_{1,n} + T_{2,n}.$$

Write

$$V(\theta) = \mathbb{E}\left\{ \frac{\text{Var}(Y_i^1 \mid X_i)}{\pi(X_i;\theta)} + \frac{\text{Var}(Y_i^0 \mid X_i)}{1-\pi(X_i;\theta)} \right\} + \text{Var}\big(\mu_1(X_i) - \mu_0(X_i)\big) \equiv V_1(\theta) + V_2.$$

Since X_i's are iid, together with $\mathbb{E}\{\mu_1(X_i) - \mu_0(X_i)\} = \psi_{\text{ATE}}$ and $\mathbb{E}\big[\{\mu_1(X_i) - \mu_0(X_i)\}^2\big] < \infty$, we have by the canonical central limit theorem (CLT) that

$$T_{2,n} \rightsquigarrow \mathcal{N}\left(0, V_2\right).$$

For $T_{1,n}$, let $H_n \equiv (X^n, A^n, Q^n)$ as in the proof of Theorem 5. We verify the conditional Lyapunov condition

$$\frac{1}{s_n^{2+\delta}} \sum_{i=1}^n \mathbb{E}(|R_{i,n}|^{2+\delta} \mid H_n) \xrightarrow{p} 0,$$

where $\delta = 2$ and $s_n^2 \equiv \sum_{i=1}^n \text{Var}(R_{i,n} \mid H_n)$, since $R_{1,n}, \ldots, R_{n,n}$ are independent conditional on H_n, as Z^n is then a deterministic function of (A^n, Q^n) under our design and the potential outcomes are independent across units.

Note that

$$s_n^2 = \sum_{i=1}^n \left\{ \frac{Z_i \text{Var}(Y_i^1 \mid X_i)}{\pi_n(X_i)^2} + \frac{(1-Z_i)\text{Var}(Y_i^0 \mid X_i)}{\{1-\pi_n(X_i)\}^2} \right\},$$

since

$$\begin{aligned}
\text{Var}(R_{i,n} \mid H_n) &= \mathbb{E}(R_{i,n}^2 \mid H_n) \\
&= \mathbb{E}\left(\left[\frac{Z_i}{\pi_n(X_i)}\{Y_i^1 - \mu_1(X_i)\} - \frac{1-Z_i}{1-\pi_n(X_i)}\{Y_i^0 - \mu_0(X_i)\} \right]^2 \,\middle|\, H_n \right) \\
&= \frac{Z_i^2}{\pi_n(X_i)^2}\mathbb{E}\left[\{Y_i^1 - \mu_1(X_i)\}^2 \mid H_n\right] + \frac{(1-Z_i)^2}{\{1-\pi_n(X_i)\}^2}\mathbb{E}\left[\{Y_i^0 - \mu_0(X_i)\}^2 \mid H_n\right] \\
&\quad - \frac{2Z_i(1-Z_i)}{\pi_n(X_i)\{1-\pi_n(X_i)\}}\mathbb{E}\left[\{Y_i^1 - \mu_1(X_i)\}\{Y_i^0 - \mu_0(X_i)\} \mid H_n\right] \\
&= \frac{Z_i \text{Var}(Y_i^1 \mid X_i)}{\pi_n(X_i)^2} + \frac{(1-Z_i)\text{Var}(Y_i^0 \mid X_i)}{\{1-\pi_n(X_i)\}^2}.
\end{aligned}$$

Now define

$$M_n \equiv \sum_{i=1}^n \mathbb{E}\big(|R_{i,n}|^4 \mid H_n\big),$$

and note that

$$\begin{aligned}
|R_{i,n}|^4 &\le 8\left| \frac{Z_i}{\pi_n(X_i)}\{Y_i^1 - \mu_1(X_i)\} \right|^4 + 8\left| \frac{1-Z_i}{1-\pi_n(X_i)}\{Y_i^0 - \mu_0(X_i)\} \right|^4 \\
&\le \frac{8}{\gamma^4}\Big(|Y_i^1 - \mu_1(X_i)|^4 + |Y_i^0 - \mu_0(X_i)|^4\Big),
\end{aligned}$$

where the first inequality follows from $|a-b|^{2+\delta} \le 2^{1+\delta}(|a|^{2+\delta} + |b|^{2+\delta})$ with $\delta = 2$, and the second from positivity. It follows that

$$\mathbb{E}\big(|R_{i,n}|^4 \mid H_n\big) \le \frac{8}{\gamma^4}\Big\{\mathbb{E}\big(|Y_i^1 - \mu_1(X_i)|^4 \mid X_i\big) + \mathbb{E}\big(|Y_i^0 - \mu_0(X_i)|^4 \mid X_i\big)\Big\},$$

and therefore

$$\frac{M_n}{n} \le \frac{8/\gamma^4}{n}\sum_{i=1}^n \Big\{\mathbb{E}\big(|Y_i^1 - \mu_1(X_i)|^4 \mid X_i\big) + \mathbb{E}\big(|Y_i^0 - \mu_0(X_i)|^4 \mid X_i\big)\Big\} = O_{\mathbb{P}}(1).$$

by the weak law of large numbers and the moment assumption.

Now decompose

$$\begin{aligned}\frac{s_n^2}{n} &= \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{Z_i \text{Var}(Y_i^1 \mid X_i)}{\pi_n(X_i)^2} + \frac{(1-Z_i)\text{Var}(Y_i^0 \mid X_i)}{\{1-\pi_n(X_i)\}^2}\right\} \\ &= \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{\text{Var}(Y_i^1 \mid X_i)}{\pi_n(X_i)} + \frac{\text{Var}(Y_i^0 \mid X_i)}{1-\pi_n(X_i)}\right\} \\ &\quad + \frac{1}{n}\sum_{i=1}^{n}\{Z_i - \pi_n(X_i)\}\left\{\frac{\text{Var}(Y_i^1 \mid X_i)}{\pi_n(X_i)^2} - \frac{\text{Var}(Y_i^0 \mid X_i)}{\{1-\pi_n(X_i)\}^2}\right\} \\ &\equiv A_n + B_n.\end{aligned}$$

We first show $A_n \xrightarrow{p} V_1(\theta)$. Let

$$a_n(X_i) \equiv \frac{\text{Var}(Y_i^1 \mid X_i)}{\pi_n(X_i)} + \frac{\text{Var}(Y_i^0 \mid X_i)}{1-\pi_n(X_i)},$$

so that

$$A_n = \frac{1}{n}\sum_{i=1}^{n} a_n(X_i).$$

Then

$$A_n - V_1(\theta) = \left[\frac{1}{n}\sum_{i=1}^{n} a_n(X_i) - \mathbb{E}\{a_n(X_i)\}\right] + \left[\mathbb{E}\{a_n(X_i)\} - V_1(\theta)\right].$$

Since $a_n(X_1), \ldots, a_n(X_n)$ are iid, positivity and the moment condition imply

$$\frac{1}{n}\sum_{i=1}^{n} a_n(X_i) \xrightarrow{p} \mathbb{E}\{a_n(X_i)\},$$

by the weak law of large numbers. Then applying Theorem 4 and dominated convergence yields

$$\mathbb{E}\{a_n(X_i)\} \to V_1(\theta),$$

and subsequently $A_n \xrightarrow{p} V_1(\theta)$.

It remains to show that $B_n = o_{\mathbb{P}}(1)$, whose proof closely follows that of $n\text{Var}(R_n) = o_{\mathbb{P}}(1)$ in Theorem 5, thanks to Lemma 4. Let

$$h_n(X_i) \equiv \frac{\text{Var}(Y_i^1 \mid X_i)}{\pi_n(X_i)^2} - \frac{\text{Var}(Y_i^0 \mid X_i)}{\{1-\pi_n(X_i)\}^2},$$

so that

$$B_n = \frac{1}{n}\sum_{i=1}^{n}\{Z_i - \pi_n(X_i)\}h_n(X_i).$$

Since $\mathbb{E}(B_n) = 0$, we have

$$n\mathbb{E}(B_n^2) = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\left[h_n(X_i)^2\{Z_i - \pi_n(X_i)\}^2\right] + \frac{2}{n}\sum_{i<j}\mathbb{E}\left[h_n(X_i)h_n(X_j)\{Z_i - \pi_n(X_i)\}\{Z_j - \pi_n(X_j)\}\right].$$

The diagonal term is again bounded as

$$\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\left[h_n(X_i)^2\{Z_i - \pi_n(X_i)\}^2\right] \le \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\{h_n(X_i)^2\} = \mathbb{E}\{h_n(X_i)^2\},$$

and the off-diagonal term, by applying Lemma 4 and $2|ab| \le a^2 + b^2$ again, as

$$\frac{2}{n}\sum_{i<j}\left|\mathbb{E}\left[h_n(X_i)h_n(X_j)\{Z_i - \pi_n(X_i)\}\{Z_j - \pi_n(X_j)\}\right]\right| \le C\mathbb{E}\{h_n(X)^2\}.$$

Now by positivity,

$$h_n(X)^2 \le \frac{2}{\gamma^4}\left[\{\mathrm{Var}(Y^1 \mid X)\}^2 + \{\mathrm{Var}(Y^0 \mid X)\}^2\right].$$

Moreover $\{\mathrm{Var}(Y^z|X)\}^2 \le \mathbb{E}\{(Y^z - \mu_z(X))^4|X\}$, so the fourth-moment assumption implies $\sup_n \mathbb{E}\{h_n(X)^2\} < \infty$. It follows that $\mathbb{E}(B_n^2) = O(n^{-1})$ and therefore $B_n = o_{\mathbb{P}}(1)$ from Chebyshev's inequality. As a result,

$$\frac{s_n^2}{n} = A_n + B_n \xrightarrow{p} V_1(\theta),$$

and

$$\frac{M_n}{s_n^4} = \frac{M_n/n}{(s_n^2/n)^2}\frac{1}{n} \xrightarrow{p} 0,$$

since $M_n/n = O_{\mathbb{P}}(1)$ and $s_n^2/n \xrightarrow{p} V_1(\theta) > 0$. Hence by Lyapunov CLT, we have

$$T_{1,n} \mid H_n \rightsquigarrow \mathcal{N}(0, V_1(\theta)).$$

It remains to show the marginal result

$$T_{1,n} \rightsquigarrow \mathcal{N}(0, V_1(\theta)),$$

and

$$T_{1,n} + T_{2,n} \rightsquigarrow \mathcal{N}\left(0, V(\theta)\right).$$

By Lévy's continuity theorem, it suffices to show pointwise convergence of the corresponding characteristic functions. Define

$$U_n \equiv \mathbb{E}(e^{itT_{1,n}} \mid H_n) - e^{-t^2 V_1(\theta)/2},$$

and note that

$$|U_n| \le \mathbb{E}(|e^{itT_{1,n}}| \mid H_n) + |e^{-t^2 V_1(\theta)/2}| \le 1 + 1 = 2.$$

Hence, combined with $\mathbb{E}\left(e^{itT_{1,n}} \mid H_n\right) \xrightarrow{p} e^{-t^2 V_1(\theta)/2}$, we have $\mathbb{E}|U_n| \to 0$. It follows that

$$\mathbb{E}(e^{itT_{1,n}}) = \mathbb{E}\big\{\mathbb{E}(e^{itT_{1,n}} \mid H_n)\big\} = e^{-t^2 V_1(\theta)/2} + \mathbb{E}(U_n),$$

and subsequently

$$\left|\mathbb{E}(e^{itT_{1,n}}) - e^{-t^2 V_1(\theta)/2}\right| = |\mathbb{E}(U_n)| \le \mathbb{E}|U_n| \to 0,$$

which implies

$$\mathbb{E}(e^{itT_{1,n}}) \to e^{-t^2 V_1(\theta)/2}.$$

Finally, note that $\mathbb{E}\big\{e^{it(T_{1,n}+T_{2,n})}\big\} = e^{t^2 V_1(\theta)/2}\mathbb{E}\big(e^{itT_{2,n}}\big) + o(1)$ since

$$\begin{aligned}
\mathbb{E}\big\{e^{it(T_{1,n}+T_{2,n})}\big\} &= \mathbb{E}\big\{e^{itT_{2,n}}\mathbb{E}(e^{itT_{1,n}} \mid H_n)\big\} \\
&= \mathbb{E}\big\{e^{itT_{2,n}}(e^{-t^2 V_1(\theta)/2} + U_n)\big\} \\
&= e^{-t^2 V_1(\theta)/2}\mathbb{E}\big(e^{itT_{2,n}}\big) + \mathbb{E}\big(e^{itT_{2,n}}U_n\big).
\end{aligned}$$

Moreover, since $|e^{itT_{2,n}}| = 1$,

$$\left|\mathbb{E}\big\{e^{it(T_{1,n}+T_{2,n})}\big\} - e^{-t^2 V_1(\theta)/2}\mathbb{E}\big(e^{itT_{2,n}}\big)\right| \le \mathbb{E}|U_n| \to 0.$$

Hence, using $\mathbb{E}(e^{itT_{2,n}}) \to e^{-t^2 V_2/2}$,

$$\mathbb{E}(e^{it(T_{1,n}+T_{2,n})}) \to e^{-t^2\{V_1(\theta)+V_2\}/2},$$

for each $t \in \mathbb{R}$, which concludes the proof. □

C.10 Proof of Corollary 1

Proof. We know from the proof of Theorem 6 that

$$\widehat{\psi}_Z - \psi_Z = \frac{\mathbb{P}_n\{f(X,Q)U\} + \mathbb{P}_n[f(X,Q)\{m(X) - \widehat{m}(X)\}]}{\mathbb{P}_n[f(X,Q)\{Z - \pi_n(X;\theta,b_n)\}]}.$$

By the same logic as in the proof of Theorem 6, $\mathbb{P}_n[f(X,Q)\{m(X) - \widehat{m}(X)\}]$ is $o_{\mathbb{P}}(n^{-1/2})$ and $\mathbb{P}_n[f(X,Q)\{Z - \pi_n(X)\}]$ converges in probability to $\mathbb{E}[f(X_i,Q_i)\{\alpha_{Q_i}(\beta,p) - \pi(X_i)\}]$. Further, since $\mathbb{E}\{f(X_i,Q_i)U_i\} = 0$ and $\mathbb{E}(U_i^4) < \infty$, applying Lyapunov CLT gives

$$\frac{1}{\sqrt{n}}\sum_{i=1}^{n} f(X_i,Q_i)U_i \rightsquigarrow \mathcal{N}\big(0, \mathbb{E}\{f(X_i,Q_i)^2\sigma(X_i)\}\big).$$

The result then follows by choosing $f(X_i,Q_i) = r_n(X_i,Q_i)/\widehat{\sigma}(X_i)$ and Slutsky's theorem. □

D Numerical experiments

All experiments were conducted on a personal laptop equipped with an Apple M4 chip and 16GB memory. All experiments completed in under 15 minutes.

D.1 Dataset

The dataset we use in our experiments is from the local government of a large U.S. county. The county employs caseworkers who connect high-risk individuals with local organizations that provide housing. As partner organization units become available, the county refers eligible individuals to the partner. Many units are reserved for specific populations, such as veterans or women experiencing domestic violence. In our simulations, we consider individuals only eligible for general-population units.

The county refers individuals to partners using first-in, first-out queues based on a computed risk score $h(X)$, which effectively determines queue assignment. The risk score aggregates the output of machine learning algorithms trained to predict several outcomes of interest, such as the risk an individual experiences homelessness in the next six months.

Individuals in lower queues have effectively no chance at treatment, so we restrict our population of interest to those with a predicted risk score $h(X)$ placing them in the top three queues. This ensures all individuals in our sample have some chance at receiving treatment under the county's current allocation process. After selecting for this population, we have 2,317 individuals in our dataset.

D.2 Synthetic DGP

In our dataset, we only observe outcomes under treatment or control, so to calculate the efficiency of our proposed estimators we generate semi-synthetic outcomes using the individual risk score $h(X)$. We assume a homogeneous treatment effect $\psi = -0.1$, corresponding to a reduction in risk. We note $|\psi| \leq \min_{x\in\mathcal{X}} h(X)$, and so all potential outcome probabilities are well-defined. Using ψ and $h(X)$, we set the potential outcomes and nuisances as

$$\begin{aligned}
\mu_0(X) &= h(X) \\
\mu_1(X) &= h(X) + \psi \\
\text{Var}(Y \mid Z = 0, X) &= \mu_0(X)(1 - \mu_0(X)) \\
\text{Var}(Y \mid Z = 1, X) &= \mu_1(X)(1 - \mu_1(X)).
\end{aligned}$$

D.3 Sample Efficiency Experiment

In this section, we provide implementation details for our two heuristic baseline algorithms and test the robustness of our results to alternative queue configurations. Similar to the experiment in the main text, we consider both a 25% and 75% budget setting, first varying the queue treatment probabilities over equally-proportioned queues, then holding treatment probabilities fixed and varying the queue assignment proportions.

D.3.1 HEURISTIC METHODS

We describe our two heuristic methods for constructing queue assignment probabilities given utility scores $u(X)$ and target queue proportions $p \in \Delta_K$.

Our first method, *Switch*, assigns individuals to queues based $u(X)$ and p, then assigns each individual a probability of switching to the adjacent higher-priority queue ($p_\uparrow$), lower-priority queue ($p_\downarrow$), or staying in their assigned queue ($1 - p_\uparrow - p_\downarrow$), returning the final assignment probabilities $\theta(X_i)$. The switching probabilities $p_\downarrow$ and$p_\uparrow$ are controlled by a single parameter $\beta \in (0, 1)$, scaled by the relative proportions of the adjacent and current queues, so that larger values of β correspond to higher switching probability.

Algorithm 1 Switch

Require: Utilities u_i, queue proportions $p \in \Delta_K$, switch probability $\beta \in (0, 1)$
Ensure: Assignment probabilities θ_{ik}

Compute cumulative proportions: $c_k \leftarrow \sum_{j \le k} p_j$
Construct bins: $b_k \leftarrow \mathrm{Quantile}(u; c_k)$ for $k = 1, \ldots, K-1$
Set bin edges: $\tau \leftarrow [-\infty, b_1, \ldots, b_{K-1}, \infty]$
Assign initial queues: $q_i \leftarrow k \quad$ such that $\quad \tau_k \le u_i < \tau_{k+1}$
for each individual i with current queue $k = q_i$ **do**
    $p_\uparrow \leftarrow \beta \cdot \dfrac{p_{k+1}}{p_{k+1} + p_k}$ if $k < K$, else 0
    $p_\downarrow \leftarrow \beta \cdot \dfrac{p_{k-1}}{p_{k-1} + p_k}$ if $k > 1$, else 0
    if $p_\uparrow + p_\downarrow > 1$ **then**
        $p_\uparrow \leftarrow \dfrac{p_\uparrow}{p_\uparrow + p_\downarrow}, \quad p_\downarrow \leftarrow \dfrac{p_\downarrow}{p_\uparrow + p_\downarrow}$
    end if
    $\theta_{i,k} \leftarrow 1 - p_\uparrow - p_\downarrow$
    $\theta_{i,k+1} \leftarrow p_\uparrow \quad$ (if $k < K$)
    $\theta_{i,k-1} \leftarrow p_\downarrow \quad$ (if $k > 1$)
end for

Our second method, *Greedy Softmax*, sequentially constructs assignment probabilities by iterating over queues in priority order. At each queue k, the algorithm computes a softmax-like score from each individual's utility u_i and residual unallocated probability mass r_i, normalizes these scores to meet the target queue assignment proportion p_k, and caps individual probabilities via an iterative rescaling procedure that ensures the budget constraint is met. The remaining mass after processing queues $1, \ldots, K-1$ is assigned to the lowest priority queue K. The scale parameter a controls how strongly utilities influence the allocation, with larger values of a concentrating mass in higher-priority queues for individuals with high utility scores.

Algorithm 2 Greedy Softmax

Require: Utilities u_i, target proportions $p \in \Delta_K$, scale a, maximum value m
Ensure: Assignment probabilities θ_{ik}

Initialize residual weights: $r_i \leftarrow 1$
Initialize $\theta_{ik} \leftarrow 0$
for $k = 1, 2, \ldots, K-1$ **do**
  Compute scores: $s_i \leftarrow \exp(a \cdot u_i \cdot r_i) - 1$
  Normalize: $s_i \leftarrow s_i \cdot \frac{p_k n}{\sum_i s_i}$
  while $\max_i s_i > m$ and $|\mathcal{C}| \cdot m < p_k n$ **do**
    $i^* \leftarrow \arg\max_i s_i$
    Add i^* to capped set $\mathcal{C}$, set $s_{i^*} \leftarrow 0$
    Rescale: $s_i \leftarrow s_i \cdot \frac{p_k n - |\mathcal{C}| \cdot m}{\sum_i s_i}$ for $i \notin \mathcal{C}$
  end while
  Set $s_i \leftarrow m$ for $i \in \mathcal{C}$
  Set $\theta_{ik} \leftarrow s_i$
  Update residuals: $r_i \leftarrow r_i - \theta_{ik}$
end for
Assign remaining mass: $\theta_{iK} \leftarrow r_i$

D.3.2 Varying Budget

We study queue configurations analogous to those in the main text under 25% and 75% budget scenarios. Figures 4 and 5 reveal qualitatively similar patterns, though the endogenous arrival estimator is less efficient due to lower variation in α_Q We examine this more directly in the next section, where we explicitly vary α_Q and study its effect on efficiency.

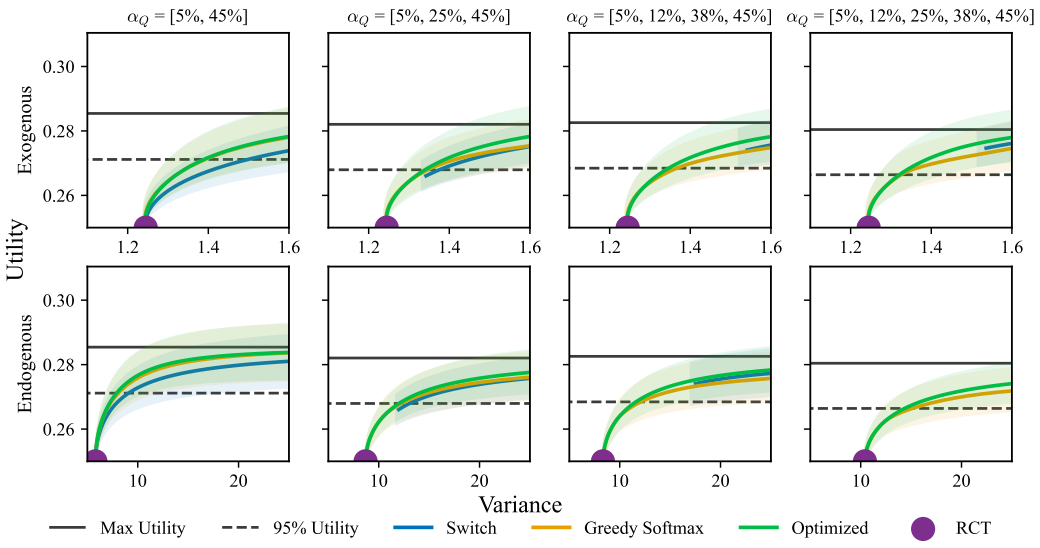

Figure 4: Pareto frontier of variance and utility under our optimized designs (green) versus two heuristic designs (orange, blue) and an RCT baseline (purple) under exogenous arrivals (top) and endogenous arrivals (bottom) across different queue configurations with a 25% treatment budget. Queues are equally sized but vary in number, with treatment probabilities shown at the top of each column.

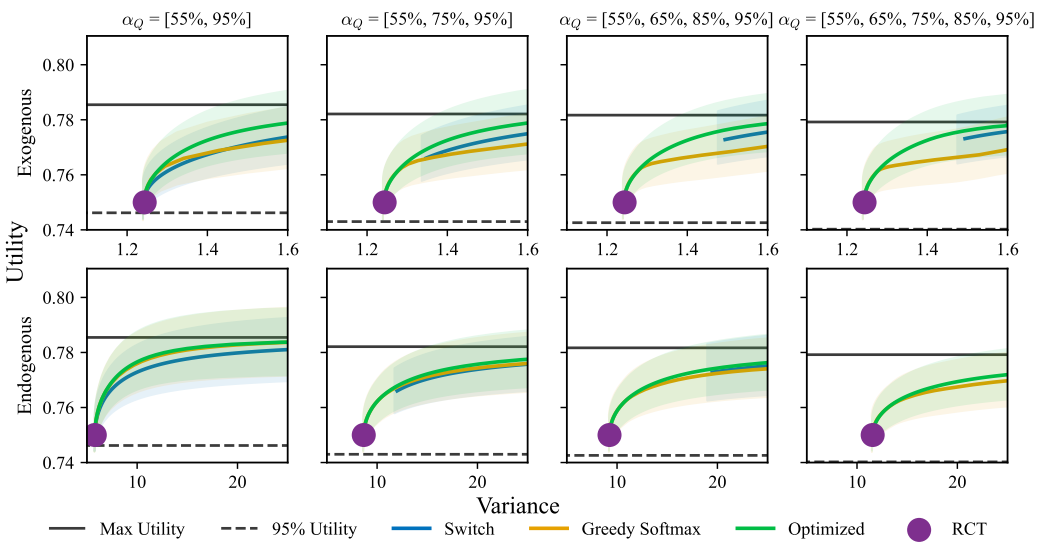

Figure 5: Pareto frontier of variance and utility under our optimized designs (green) versus two heuristic designs (orange, blue) and an RCT baseline (purple) under exogenous arrivals (top) and endogenous arrivals (bottom) across different queue configurations with a 50% treatment budget. Queues are equally sized but vary in number, with treatment probabilities shown at the top of each column. All designs, including the RCT, achieve more than 95% utility.

D.3.3 Variance of α_Q on Efficiency

Figure 6 shows that the between-queue variance of treatment probabilities greatly influences the efficiency of the endogenous estimator. As the variance of α_Q increases, the endogenous estimator becomes more efficient, and at extreme treatment probabilities, 1% in the lowest-priority queue and 99% in the highest-priority queue, it achieves efficiency comparable to the exogenous estimator.

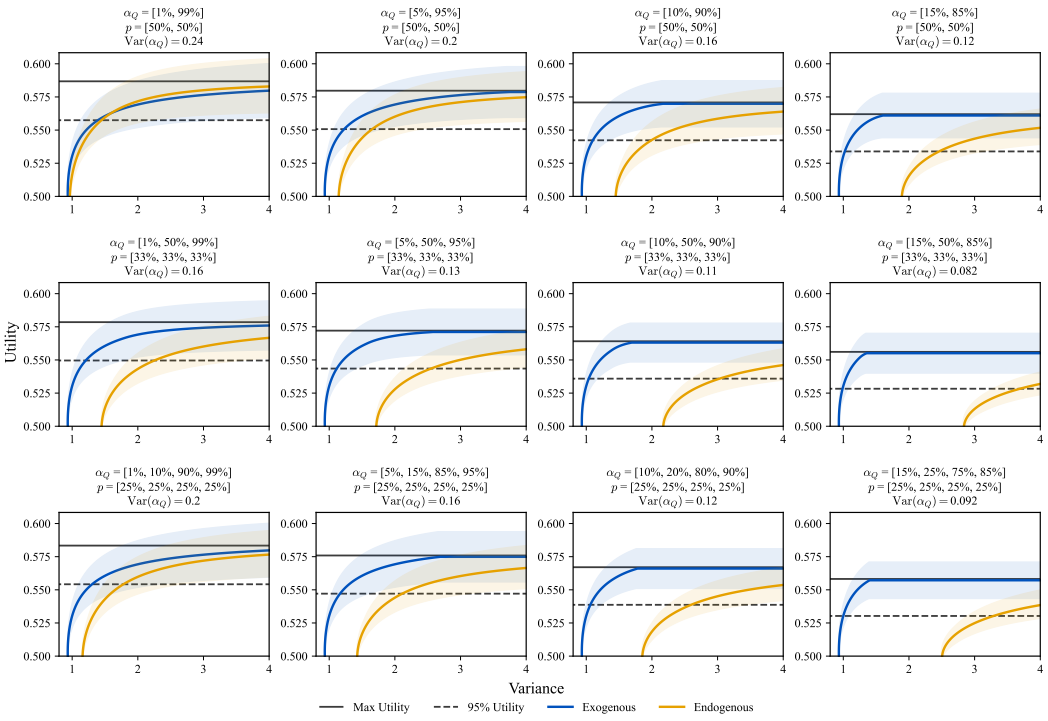

Figure 6: Pareto frontier of variance and utility under the endogenous design (orange) and exogenous design (blue) across different queue configurations with a 50% treatment budget. Treatment probabilities and queue assignment proportions are shown at the top of each column.

D.3.4 Varying Queue Assignment Proportions

We vary the queue proportions while keeping the budget fixed, allocating more individuals to lower-priority queues. As with the varying budget setting, Figure 7 reveals that the optimized designs remain competitive across all queue configurations, while neither heuristic dominates the other.

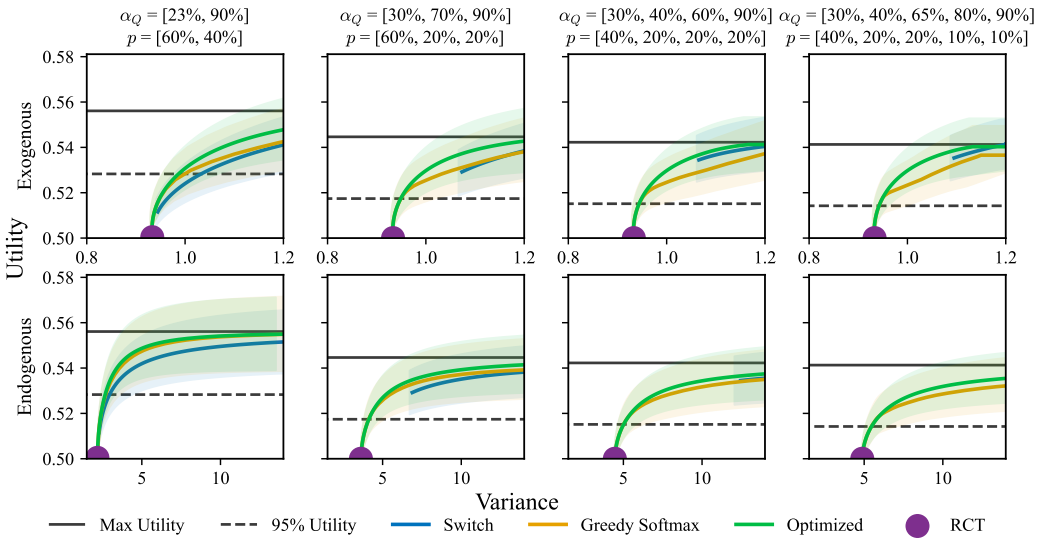

Figure 7: Pareto frontier of variance and utility under our optimized designs (green) versus two heuristic designs (orange, blue) and an RCT baseline (purple) under exogenous arrivals (top) and endogenous arrivals (bottom) across different queue configurations with a 50% treatment budget. Treatment probabilities and queue assignment proportions are shown at the top of each column.

D.4 Bias Experiment

Potential outcomes are generated according to the partially linear model

$$Y^1 = \psi + h(X) + U,$$
$$Y^0 = h(X) + U,$$

where $\psi = -0.1$ is the homogeneous treatment effect and

$$U \sim \text{Uniform}(-0.2h(X),\ 0.2h(X)),$$

The arrival A is set as the indices of U sorted from highest to lowest.

For both the endogenous and exogenous settings, we use A, Y^1, Y^0, and the corresponding $\theta(X)$, to simulate observed treatment assignments Z and outcomes Y. To study bias, we also calculate the misspecified nuisance functions $E[Y^1 \mid X]$ and $E[Y^0 \mid X]$ that omit U, introducing misspecification into the exogenous estimator ψ_{DR} that relies on these models. The realized outcomes, treatments, and known nuisance functions are then plugged into their corresponding estimators ψ_{DR} in the exogenous setting and ψ_Z in the endogenous setting. We repeat this procedure 10,000 times and report the average bias $\mathbb{E}[\widehat{\psi}.] - \psi$ for both estimators.

E Public-Data Replication

To make our experimental comparisons reproducible without access to the administrative data, we additionally evaluate our designs using publicly available data from the Medical Expenditure Panel Survey (MEPS). Specifically, we use the Panel 21 longitudinal file, HC-202, which contains observations from 2016 and 2017.[2] This experiment is aimed at providing a public-data replication of the experimental framework in Section 5: we construct predictive risk scores from the public data, then use these scores to evaluate sample efficiency and robustness to arrival endogeneity under synthetic potential outcomes.

[2] https://meps.ahrq.gov/mepsweb/data_stats/download_data_files_detail.jsp?cboPufNumber=HC-202

Dataset and risk scores: We restrict the sample to individuals aged 18–64 in 2016 who participated in all five survey rounds and have an observed inpatient discharge count in 2017; this yields 8,774 individuals. We then split this sample by household into training and evaluation data containing 5,266 and 3,508 individuals, respectively.

Using the training sample, we fit an unregularized logistic regression to predict whether an individual has at least one inpatient discharge in 2017 based on the covariates: age, sex, income category, insurance coverage, inpatient discharge and emergency department visit counts in 2016, and perceived health status in the first survey round. Missing numeric and categorical predictors are imputed using training-sample medians and modes, respectively. We set $u(X) = h(X)$ throughout, where $h(X)$ is the fitted risk score. We use the evaluation sample as a fixed, unweighted empirical population, holding its risk scores fixed across simulations.

Synthetic DGP: As in D.4, potential outcomes are generated according to the partially linear model

$$Y^1 = \psi + h(X) + U,$$
$$Y^0 = h(X) + U,$$

where $\psi = -0.05$ is the homogeneous treatment effect and

$$U \sim \text{Uniform}(-0.1,\ 0.1)$$

independently of X. The simulated outcomes are continuous and are not truncated to $[0, 1]$.

For exogenous arrivals, we draw $A \sim \text{Uniform}(0, 1)$ independently of (X, U). For endogenous arrivals, we instead set

$$A = \frac{0.1 - U}{0.2},$$

so that individuals with larger values of the unobserved outcome component arrive earlier. We implement the prescribed queue treatment probabilities through arrival cutoffs $Z = \mathbb{1}\{A \leq \alpha_Q\}$; thus, $\mathbb{P}(Z = 1 \mid X, Q = k) = \alpha_k$ in both settings.

Sample efficiency. We use the four queue configurations in Figure 1, with equal queue proportions and a 50% expected treatment budget. We retain the optimization objectives and negative-entropy regularization coefficient of one used in our implementation. For each configuration, we solve the optimization problem at 20 equally spaced utility constraints from 0.5 to 99% of the maximum achievable utility.

We compare the optimized designs to *Switch*, *Greedy Softmax*, and uniform queue assignment. For Switch, we use 100 equally spaced values of β from 0.02 to 1; for Greedy Softmax, we use 100 logarithmically spaced values of a from 10^{-2} to 10^3. We split probability mass at quantile boundaries to enforce the specified queue proportions and cap each Greedy Softmax allocation at the individual's remaining probability mass.

Figure 8 reveals a similar trade-off between utility and statistical efficiency to the housing allocation experiment. The endogenous design generally has higher variance than the exogenous design, with a larger increase in variance as utility approaches its maximum. The optimized designs remain competitive with the heuristic methods across queue configurations. With more queues, Switch covers a restricted range of utility values, while the optimized designs span the trade-off from uniform queue assignment to near-maximal utility.

Arrival endogeneity. We repeat the bias experiment using two equally sized queues with treatment probabilities $\{1 - \gamma, \gamma\}$, where γ takes five equally spaced values from 0.6 to 0.99. For each configuration, we solve for both designs over the same 20-point utility grid rule and average the estimation error over 10,000 simulations. We use the known values of $\pi(X)$, $m(X) = h(X) + \psi\pi(X)$, and the noise variance $0.1^2/3$ for the endogenous estimator. For the exogenous estimator, we use $\mu_0(X) = h(X)$ and $\mu_1(X) = h(X) + \psi$, which are the potential-outcome means but omit the selection on ϵ induced by endogenous arrivals. As in Figure 2, Figure 9 shows substantial bias for the exogenous estimator and much smaller bias for the endogenous estimator.

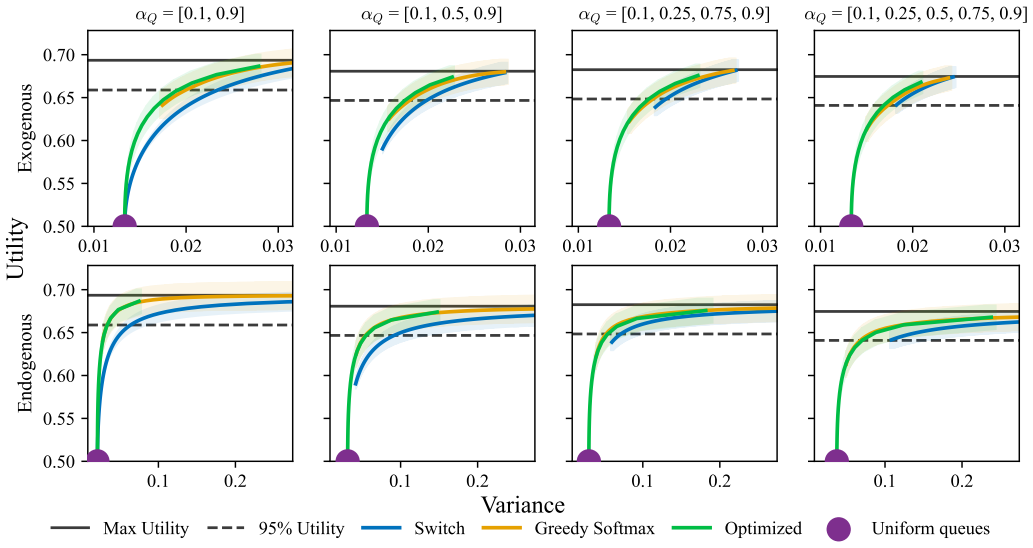

Figure 8: Pareto frontier of variance and utility under our optimized designs (green) versus two heuristic designs (orange, blue) and uniform queue assignment (purple) using MEPS data, under exogenous arrivals (top) and endogenous arrivals (bottom) across different queue configurations with a 50% expected treatment budget. As before, exogenous designs in general are more efficient than endogenous designs.

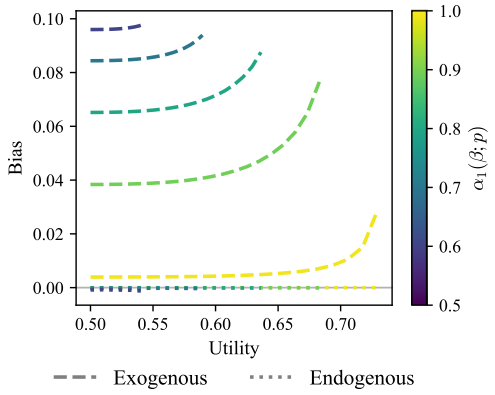

Figure 9: Bias under arrival endogeneity using MEPS data for two queues with varying treatment probabilities.