

Discussion Paper No. 1308

ISSN (Print) 0473-453X

ISSN (Online) 2435-0982

**OPTIMAL TREATMENT ASSIGNMENT RULES
UNDER CAPACITY CONSTRAINTS**

Keita Sunada
Kohei Izumi

March 2026

The Institute of Social and Economic Research
The University of Osaka
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

OPTIMAL TREATMENT ASSIGNMENT RULES UNDER CAPACITY CONSTRAINTS

KEITA SUNADA AND KOHEI IZUMI

ABSTRACT. We study treatment assignment when treatments are limited in supply, where a planner aims to maximize social welfare by assigning treatments based on observable covariates. Such constraints are common when treatments are scarce and costly, but they complicate the analysis of optimal assignment rules because assignment probabilities must be coordinated across the entire covariate distribution. We develop a new approach that reformulates the planner’s problem as an *optimal transport* problem, which makes the constraints analytically tractable. Using a limits of experiments framework, we establish local asymptotic optimality results for two canonical decision rules—the plug-in rule and the Bayesian rule. We show that the former rule can dominate the latter rule, with simulations demonstrating sizable risk reductions. An empirical illustration using school voucher program data from Angrist *et al.* (2006) demonstrates how the two rules differ in practice.

JEL Classification: C10, C18, C21, C44, C61

Keywords: Treatment assignment, capacity constraints, optimal transport, statistical decision theory, external validity.

1. INTRODUCTION

When a social planner allocates treatments such as school vouchers or subsidies, an important question arises: who should be treated given available data? The planner typically does not know the true treatment effects but can rely on observed data from experimental or observational studies. The literature on treatment assignment problems has examined how to use available data to design treatment assignment rules that maximize social welfare (e.g., Manski, 2004; Stoye, 2009; Hirano and Porter, 2009). While this literature has made substantial progress, much less is known about how to design optimal rules when treatments are limited in supply. In many real-world settings, treatment assignment is subject to such constraints—budgets, supplies, or available slots are limited, so not everyone who could benefit can be treated.

This paper extends Hirano and Porter (2009) to derive optimal rules when treatments are limited in supply. To accommodate these constraints, we reformulate the planner’s problem as an *optimal transport* problem, and establish locally asymptotically optimal assignment rules—results that, to

Date: March 26, 2026 (First version: June 13, 2025).

We are particularly grateful to Nese Yildiz for her guidance and support. We thank Bin Chen, Kazuma Inagaki, Toru Kitagawa, Laura Liu, Yukun Ma, Eric Mbakop, Tai Otsu, Yuya Sasaki, John Singleton, Takeaki Sunada, Takuro Yamashita, and seminar participants at Midwest Econometrics Study Group (UIUC), Canadian Econometrics Study Group in Ottawa, the University of Iowa, and Aarhus University for their helpful comments and suggestions. This paper was awarded the 2025 Moriguchi Prize by the Institute of Social and Economic Research, the University of Osaka.

our knowledge, have not yet been formally derived in the literature. The framework also naturally extends to settings with non-binary treatments.

At first glance, one might think that the constrained problem is only a straightforward extension of the setting without constraints on available treatments—simply add a quota. Since Hirano and Porter (2009) show that assigning treatment to individuals with (efficiently) estimated positive effects is locally asymptotically optimal when there are no constraints on available treatments, it is tempting to conclude that ranking individuals by their estimated effects and treating them in descending order until the quota is filled would also be optimal. While this intuition turns out to be correct under suitable smoothness conditions on the planner’s utility function—a result we establish formally—it is far from straightforward to prove. The presence of capacity constraints makes the derivation of optimal assignment rules substantially more complex, as explained below.

We consider a planner who aims to maximize social welfare by designing treatment rules based on observable covariates, subject to a capacity constraint. Let $w(\theta, x, t)$ denote the planner’s utility from assigning treatment $T = t$ to an individual with covariate $X = x$, where $\theta \in \mathbb{R}^k$ characterizes the treatment effects. For example, if $w(\theta, x, t)$ is the conditional mean of the potential outcome given (x, t) , then θ is a parameter indexing the conditional potential outcome distribution. Suppose that the true parameter θ_0 can be efficiently estimated by an estimator $\hat{\theta}_n$ constructed from experimental or observational data. For simplicity, consider a binary treatment setting where a fixed fraction p of the population is to be treated. When θ_0 is known, the planner’s problem is

$$\max_{\delta} \int_{\mathcal{X}} \{w(\theta_0, x, 1)\delta(x) + w(\theta_0, x, 0)(1 - \delta(x))\} dF_X(x), \quad \text{s.t.} \quad \int_{\mathcal{X}} \delta(x) dF_X(x) = p,$$

where F_X is the distribution of covariates and $\delta(x)$ is the probability of assigning treatment to $X = x$. In reality, θ_0 is not known to the planner. Thus the planner utilizes the available data to design treatment assignment rules.

Under the current setup, however, we cannot simply follow a standard approach to analyzing the asymptotically optimal data-driven rules taken by Hirano and Porter (2009) and others. This approach first derives the optimal rule *for each covariate* x in an asymptotically equivalent (and simpler) problem. It then constructs a sequence of feasible data-driven rules that asymptotically match the optimal rule in the limit. The key is that the limiting problem becomes tractable because the optimal rule in the limit is selected from the set of rules obtained as *asymptotic representations* of data-driven rules, which only depend on random variables following a shifted normal and a uniform distribution (van der Vaart, 1991, Theorem 3.1).

Under capacity constraints, this pointwise approach fails because the planner must coordinate assignment probabilities across the entire covariate distribution F_X . A natural alternative is to define the planner’s action space as the set of (measurable) functions

$$\mathcal{F} := \left\{ \delta : \mathcal{X} \rightarrow [0, 1] \mid \int_{\mathcal{X}} \delta(x) dF_X(x) = p \right\}$$

and work directly with \mathcal{F} . When θ_0 is known, the constrained optimization over \mathcal{F} can be solved via a Lagrangian, reducing the problem to an unconstrained one. However, this approach encounters two difficulties in the statistical problem where θ_0 is unknown. First, the Lagrange multiplier depends on θ , so the Lagrangian reduction does not directly apply when θ_0 is unknown and must be estimated from data—the multiplier must be jointly determined with the parameter, complicating the construction and analysis of data-driven rules. Second, \mathcal{F} fails to be compact under common norms when the support of X is not finite or countable, making it difficult to obtain the asymptotic representations of assignment rules.¹ To proceed, we need a formulation that circumvents these difficulties.

We address these difficulties by reformulating the planner’s constrained problem as an optimal transport problem. The key observation is that, when the capacity constraint binds exactly, the (unconditional) distribution of treatment assignments must be a Bernoulli distribution with success probability p . The planner’s problem can then be framed as transporting the mass of F_X into the Bernoulli distribution F_T , in a way that maximizes the social welfare. This reformulation allows us to treat the action space as a set of *couplings*—joint distributions on $\mathcal{X} \times \{0, 1\}$ whose marginals are F_X and F_T —equipped with the Wasserstein distance. This new action space is a compact and convex metric space, which resolves both difficulties above: compactness enables the asymptotic representations, and the capacity constraint is automatically satisfied by any coupling in \mathcal{M} , eliminating the need for a Lagrangian. Moreover, this framework naturally extends beyond binary treatments to general discrete or continuous settings, where the Lagrangian approach becomes increasingly intractable. To our knowledge, this is the first study to reformulate the optimal treatment assignment problem into an optimal transport problem.

This reformulation enables us to establish decision-theoretic optimality results for treatment assignment rules under capacity constraints. The existing local asymptotic frameworks for treatment choice, including Hirano and Porter (2009) and Christensen *et al.* (2026), do not directly accommodate capacity constraints. Our OT formulation resolves this, but in doing so introduces a new technical challenge: the action space \mathcal{M} of couplings is infinite-dimensional, unlike the finite action spaces considered in the prior literature. We employ posterior-risk minimization in the Gaussian limit experiment as our optimality criterion. This criterion is essentially equivalent to the standard average optimality used in Hirano and Porter (2009) and Christensen *et al.* (2026) whenever the average risk is well-defined, while remaining applicable even in settings where average optimality may not be.

We analyze two canonical assignment rules: the plug-in rule, which replaces θ_0 with its efficient estimator $\hat{\theta}_n$, and the Bayesian rule, which forms the posterior on the parameter space Θ based on

¹By van der Vaart (1991, Theorem 3.1), \mathcal{F} needs to be complete and separable to obtain the asymptotic representations of assignment rules. Under the sup-norm, \mathcal{F} is typically not separable when \mathcal{X} is not countable. Under L_p -norm; i.e., $\|\cdot\|_p = \left(\int |\cdot|^p dF_X(x)\right)^{1/p}$, \mathcal{F} is separable and complete, but not compact unless \mathcal{X} is finite. Thus it is not straightforward to find a suitable norm to work with \mathcal{F} .

the observed data. Our main results are as follows: (i) both rules are optimal in the limit experiment under capacity constraints when the planner’s utility function is continuously differentiable, (ii) the plug-in rule is generally suboptimal, while the Bayesian rule remains optimal when the planner’s utility function is only directionally differentiable.

Directionally differentiable utility functions often arise in empirically relevant settings. For instance, they arise when the (conditional) potential outcome distributions may differ slightly between the target and training populations, and the planner adopts a maximin-type utility function that is robust to such distributional shifts (Adjaho and Christensen, 2023).

Our results are consistent with the existing literature. The plug-in rule is known to be average optimal under point-identified models when the planner’s utility function is continuously differentiable (Hirano and Porter, 2009), while the Bayesian rule is known to be average optimal under partially identified models when the planner’s utility function is directionally differentiable (Christensen *et al.*, 2026), where the lack of point identification similarly precludes full differentiability and renders the plug-in rule suboptimal. While our analysis concerns a different environment by featuring capacity constraints, our limit-experiment results imply the corresponding average-optimality conclusions within our model whenever average optimality is well-defined. Thus our findings complement the existing literature by establishing analogous optimality properties in capacity-constrained settings. On the technical side, our proof strategy avoids imposing any uniqueness condition on the maximizer of the welfare function. In the coupling space \mathcal{M} , the welfare-maximizing set A_0 is generically non-singleton, which precludes the direct application of standard argmax theorems for establishing the weak convergence of the Bayesian rule. We circumvent this by defining the Bayesian rule as the unique minimizer of a strictly convex penalty function over the set of welfare maximizers. This construction enables the weak convergence analysis without requiring the “no first-order ties” condition of Christensen *et al.* (2026) or the uniqueness assumption of Xu (2024).

To quantify the performance gap between the two rules, we conduct a simulation study. The results support our theoretical findings. Specifically: (i) the Bayesian rule achieves a substantially lower risk than the plug-in rule in small samples for both continuously and directionally differentiable utility functions, (ii) the two rules behave similarly under the continuously differentiable utility function in larger samples, and (iii) the Bayesian rule continues to outperform the plug-in rule under the directionally differentiable utility function in larger samples.

We illustrate our methods using data from Angrist *et al.* (2006), who study the impact of receiving a randomly assigned voucher—allowing students to attend private high schools—on educational attainment seven years later. We hypothetically treat the marginal distribution of covariates (age and sex) in the observed sample as that of the target population, and compute both the plug-in rule and the Bayesian rule. The two rules yield identical allocations under a continuously differentiable utility function but differ under the maximin-type directionally differentiable utility function mentioned above.

1.1. Related literature. This study builds on the literature on the statistical treatment assignment problems in econometrics, where the pioneering works include Manski (2004) and Dehejia (2005). Within this expanding literature, our main contribution is to provide a decision-theoretic optimality result under capacity constraints on available treatments. Previous studies have established the decision-theoretic optimality of treatment rules in several settings including: (i) point-identified smooth (semi-)parametric models under local asymptotics (Hirano and Porter, 2009; Masten, 2023), (ii) partially-identified smooth (semi-)parametric models under local asymptotics (Christensen *et al.*, 2026; Kido, 2023; Xu, 2024), (iii) point-identified models with finite samples (Stoye, 2009; Stoye, 2012; Tetenov, 2012; Guggenberger *et al.*, 2024; Kitagawa *et al.*, 2024; Chen and Guggenberger, 2025), (iv) partially-identified models with finite samples (Manski, 2007; Stoye, 2012; Yata, 2023; Ishihara and Kitagawa, 2024; Aradillas Fernández *et al.*, 2024; Montiel Olea *et al.*, 2026). However, none of these studies consider the capacity constraints in the way we do. To the best of our knowledge, this is the first study to establish an optimality result under such constraints within the framework of Hirano and Porter (2009), extended to cover non-binary treatments—both discrete and continuous.

Besides Hirano and Porter (2009), the most closely related paper is Christensen *et al.* (2026). They extend Hirano and Porter (2009) to allow for partially identified parameters and (non-randomized) discrete actions, and show the asymptotic optimality of the Bayesian rule. Our setting differs in that the action space is the set of couplings of F_X and F_T (see notation below) equipped with the Wasserstein distance, which naturally allows for randomized assignment rules. This non-standard action space requires nontrivial extensions to analyze the asymptotic properties of the Bayesian rule, which we address by drawing on tools from optimal transport. In contrast to their framework, we focus on point-identified parameters. Xu (2024) further extends this line of work to continuous decision problems (point decisions) using an expansion-based approach, accommodating non-smooth loss functions in finite-dimensional action spaces. Our setting differs in that the action space is the infinite-dimensional space of couplings arising from the OT formulation of the capacity-constrained problem.

There also exist studies that incorporate exogenously given constraints into the treatment assignment problem. Bhattacharya and Dupas (2012) impose the capacity constraints in the way we do, but focus on the estimation and inference of a nonparametric plug-in rule. Other papers adopting the *empirical welfare maximization* approach allow for various types of constraints, including the capacity constraints (Kitagawa and Tetenov, 2018; Athey and Wager, 2021; Mbakop and Tabord-Meehan, 2021; Sun, 2024). Kitagawa and Tetenov (2018) show the optimality of their proposed rule in terms of the welfare convergence rate, which measures how quickly the average welfare achieved by the proposed rule converges to the maximum welfare under the true data generating process. In our semiparametric framework with capacity constraints, the empirical welfare maximization rule of Kitagawa and Tetenov (2018) reduces to the plug-in rule that maximizes $W(\hat{\theta}_n, \mu)$ over couplings. Our results show that this rule is optimal in the limit experiment when the planner’s utility

w is continuously differentiable, but can be suboptimal when w is only directionally differentiable (Section 3.4). In the latter case, the Bayesian rule strictly dominates—a finding that the finite-sample minimax regret framework of Kitagawa and Tetenov (2018) does not capture, as it does not distinguish between smooth and non-smooth welfare functions in terms of optimality.

Some recent works utilize tools from optimal transport theory in the literature of treatment assignment problems. Kido (2022) and Adjaho and Christensen (2023) study the external validity of treatment choices by measuring the difference in potential outcome distributions between the training and target populations using the Wasserstein distance. Hazard and Kitagawa (2025) formulate a learning problem of optimal matching policies in a two-sided market as an empirical optimal transport problem, and derive a welfare regret bound for their estimated policy.

Our work also relates to the growing field of statistical methods for optimal transport problems (Chewi *et al.*, 2024), as we study the local asymptotic properties of transport maps of an optimal transport problem where the cost function is indexed by parameters that can be efficiently estimated.

1.2. Structure of the paper. The remainder of the paper is organized as follows. Section 2 formulates the planner’s problem and introduces the data generating process. Section 3 introduces the decision theoretic framework and define the plug-in rule and the Bayesian rule. Then the optimality results are stated. Section 4 provides a simulation study to evaluate the finite sample performance of rules. Section 5 illustrates our methods using the data from Angrist *et al.* (2006). Finally, Section 6 concludes. All of the proofs are relegated to Appendix.

1.3. Notation. A function $f : \Theta \subset \mathbb{R}^k \rightarrow \mathbb{R}$ is (*Hadamard*) *directionally differentiable* at θ_0 if there is a continuous function $\dot{f}_{\theta_0} : \mathbb{R}^k \rightarrow \mathbb{R}$ such that

$$\lim_{n \rightarrow \infty} \left| \frac{f(\theta_0 + t_n h_n) - f(\theta_0)}{t_n} - \dot{f}_{\theta_0}(h) \right| = 0$$

for all sequences $\{t_n\} \subset \mathbb{R}_+$ and $\{h_n\} \subset \mathbb{R}^k$ such that $t_n \downarrow 0$, $h_n \rightarrow h \in \mathbb{R}^k$ as $n \rightarrow \infty$ and $\theta_0 + t_n h_n \in \Theta$ for all n . It is worth noting that this requires \dot{f}_{θ_0} need to be continuous, but not to be linear.

Let P and Q be Borel probability measures on \mathcal{A} and \mathcal{B} , respectively. A joint distribution μ on $\mathcal{A} \times \mathcal{B}$ is called a *coupling* of P and Q if its marginals are P and Q ; that is, $\mu(\mathcal{A} \times \mathcal{B}) = P(\mathcal{A})$ and $\mu(\mathcal{A} \times \mathcal{B}) = Q(\mathcal{B})$ for any measurable sets A and B .

2. SETTING

The setting of this paper closely follows that of Hirano and Porter (2009). We consider a social planner who assigns a treatment T to individuals based on their observable covariates X . Let F_X denote the marginal distribution of X in the target population, with support \mathcal{X} . We assume that F_X is known to the planner. The planner can fractionally (probabilistically) assign treatment $T = t$ to an individual with covariate $X = x$.

Let $Y(t)$ denote potential outcomes under treatment $T = t$. In contrast to Hirano and Porter (2009), we distinguish between the conditional potential outcome distribution in the target population and in the training population. We denote the conditional distribution of $Y(t)$ in the training population as $F_t(\cdot|x, \theta)$, where $F_t(\cdot|x, \theta)$ belongs to families of distributions indexed by a parameter $\theta \in \Theta \subset \mathbb{R}^k$. The planner must learn θ from the available data from experimental or observational studies.

2.1. The planner's preferences. The planner's utility for assigning treatment $T = t$ to an individual with covariate $X = x$ depends on the conditional distribution $F_t(\cdot|x, \theta)$ via a functional w . For the shorthand notation, we write

$$w(\theta, x, t) := w(F_t(\cdot|x, \theta)).$$

We consider two scenarios. First, $w(\theta, x, t)$ is fully differentiable in θ . Second, $w(\theta, x, t)$ is only directionally differentiable in θ . Two examples corresponding to each scenario are given as follows.

Example 2.1. When the planner is interested in the (conditional) mean outcome, then a natural choice of utility function is

$$w(\theta, x, t) = \int y dF_t(y|x, \theta).$$

This choice is standard in the literature and appropriate especially when the target and training populations are assumed to have the same conditional potential outcome distribution. ■

Example 2.2. When the conditional potential outcome distributions may differ slightly across the target and training populations, the planner may wish to adopt a utility function that is robust to distributional shifts. Following Adjaho and Christensen (2023), we formalize such a utility function.² We first define an ε -neighborhood of $F_t(\cdot|x, \theta)$ as

$$\mathcal{N}_\varepsilon := \left\{ G_t(\cdot|x) : \tilde{d}_W(G_t(\cdot|x), F_t(\cdot|x, \theta)) \leq \varepsilon \right\},$$

where $\varepsilon > 0$ is a measure of neighborhood size, and \tilde{d}_W is the Wasserstein distance of order 1.³ Given $\lambda \in [0, 1]$, the planner's utility for assigning treatment $T = t$ for individuals with covariate $X = x$ is then defined as

$$w(\theta, x, t) = \lambda \int y dF_t(y|x, \theta) + (1 - \lambda) \inf_{G_t(\cdot|x) \in \mathcal{N}_\varepsilon} \int y dG_t(y|x).$$

This formulation corresponds to maxmin preferences of an ambiguity-averse decision maker (Gilboa and Schmeidler, 1989). In the second term, since the true target distributions are unknown, the

²For alternative approaches to robust welfare, see Si *et al.* (2020), Kido (2022), and Qi *et al.* (2023).

³Formally,

$$\tilde{d}_W(F_t(\cdot|x), G_t(\cdot|x, \theta)) := \inf_{\pi \in \Pi(F_t, G_t)} \int |y - \tilde{y}| d\pi(y, \tilde{y}),$$

where $\Pi(F_t, G_t)$ denotes all couplings of F_t and G_t . Note that \tilde{d}_W differs from d_W defined in (2.3).

planner computes the conditional mean under the worst-case distribution within the neighborhood of $F_t(\cdot|x, \theta)$, treating it as a fixed reference prior on $Y(t)$ given x .⁴

By Adjaho and Christensen (2023, Remark 2.2), this utility function can be rewritten as

$$w(\theta, x, t) = \lambda \int y dF_t(y|x, \theta) + (1 - \lambda) \max \left\{ \int y dF_t(y|x, \theta) - \varepsilon, y_\ell(t) \right\}, \quad (2.1)$$

where $y_\ell(t)$ is the possible minimum values in the support of $Y(t)$. From this expression, one finds that (2.1) is only directionally differentiable, and the directional derivative of the second term of (2.1) for direction h at θ is given by $(1 - \lambda)$ times

$$\begin{cases} \left(\frac{\partial}{\partial \theta} \int y dF_t(y|x, \theta) \right)^\top h & \text{if } \int y dF_t(y|x, \theta) - \varepsilon > y_\ell(t), \\ \max \left\{ \left(\frac{\partial}{\partial \theta} \int y dF_t(y|x, \theta) \right)^\top h, 0 \right\} & \text{if } \int y dF_t(y|x, \theta) - \varepsilon = y_\ell(t), \\ 0 & \text{if } \int y dF_t(y|x, \theta) - \varepsilon < y_\ell(t). \end{cases}$$

We note that the non-linearity of this directional derivative corresponds to the failure of the full differentiability of (2.1) at θ . ■

Remark 2.3. There are two distinct sources of non-differentiability in this problem. The first arises from the OT structure: the value function $W_{\mathcal{M}}^*(\theta) = \max_{\mu \in \mathcal{M}} W(\theta, \mu)$ is only Hadamard directionally differentiable, even when $w(\theta, x, t)$ is smooth (Lemma B.1). This is a generic consequence of the pointwise maximum of a family of functions. The second source is the non-differentiability of $w(\theta, x, t)$ itself, as illustrated in Example 2.2 above. For the optimality results developed in Section 3, it is the second source that is critical: when w is continuously differentiable, both the plug-in and Bayesian rules are optimal in the limit experiment, despite the non-differentiability of $W_{\mathcal{M}}^*$. By contrast, when w is only directionally differentiable, the plug-in rule may fail to be optimal, while the Bayesian rule remains optimal.

2.2. The planner's problem as optimal transport. We consider a setting where the planner faces the capacity constraints on the available treatments. To illustrate, consider a simple binary treatment case, $T \in \mathcal{T} := \{0, 1\}$, where a fraction p of the target population is to be treated. We assume that the capacity constraint binds exactly. Let $\delta(x)$ denote the probability of assigning treatment $T = t$ to individuals with $X = x$. Then, under the capacity constraint, the planner's problem can be written as

$$\max_{\delta \in \mathcal{F}} \int_{\mathcal{X}} \{w(\theta, x, 1)\delta(x) + w(\theta, x, 0)(1 - \delta(x))\} dF_X(x), \quad (2.2)$$

We now show how to convert this constrained optimization problem into more tractable one. Observe that the distribution of treatment assignments must be a Bernoulli distribution with the success probability p . With this in mind, the planner's problem can be seen as an optimal transport

⁴Several recent studies incorporate non-Bayesian preferences that arise naturally in settings with set-identifiable parameters (Giacomini and Kitagawa, 2021; Aradillas Fernández *et al.*, 2024; Christensen *et al.*, 2026). Banerjee *et al.* (2020) adopt maxmin preferences to study of the optimal experimental design problems.

problem: the planner transports the mass of F_X into F_T , the Bernoulli distribution, in a way that maximizes the social welfare.

Formally, let \mathcal{M}_a be the set of all couplings of F_X and F_T . Let d be a distance function on $\mathcal{X} \times \mathcal{T}$, and define the Wasserstein distance of order 1 as

$$d_W(\mu, \nu) := \inf_{\gamma \in \Gamma(\mu, \nu)} \int d((x, t), (x', t')) d\gamma((x, t), (x', t')), \quad (2.3)$$

where $\Gamma(\mu, \nu)$ is the set of couplings whose marginals are μ and ν . We focus on couplings that have a finite first moment:

$$\mathcal{M} := \left\{ \mu \in \mathcal{M}_a : \int d((x_0, t_0), (x, t)) d\mu < +\infty \right\},$$

for some arbitrary $(x_0, t_0) \in \mathcal{X} \times \mathcal{T}$. Then (\mathcal{M}, d_W) becomes a metric space, and d_W is finite on \mathcal{M} (Villani, 2009, Theorem 6.9). Using this setup, the original constrained problem (2.2) is equivalent to the following:

$$\max_{\mu \in \mathcal{M}} W(\theta, \mu), \quad (2.4)$$

where

$$W(\theta, \mu) := \int_{\mathcal{X} \times \mathcal{T}} w(\theta, x, t) d\mu = \int_{\mathcal{X}} \int_{\mathcal{T}} w(\theta, x, t) d\mu(t|x) dF_X(x).$$

Hence, the action space of the planner is the space of couplings (\mathcal{M}, d_W) .

There are three important remarks regarding this optimal transport formulation. First, the capacity constraint is automatically satisfied by any coupling in \mathcal{M} , making the problem effectively unconstrained. Second, this reformulation is computationally attractive, as one can use an existing software for optimal transport. Finally, this approach can easily accommodate non-binary treatment settings. We assume that T follows a distribution F_T , determined by the capacity constraints, with support \mathcal{T} . We remark that $d\mu(t|x)$ becomes a conditional probability measure when F_T is continuous.

Remark 2.4. The capacity constraint is satisfied at the population level by any coupling $\mu \in \mathcal{M}$. However, \mathcal{M} contains stochastic allocations. For example, under some $\mu \in \mathcal{M}$, an individual may be assigned to t with probability 1/2 and to $t' \neq t$ with probability 1/2. In this sense, our rule may violate the capacity constraint ex-post. But the concern about stochastic allocations becomes less serious when many individuals share the same value of X . In that case, a stochastic allocation at a given X can be interpreted as a fractional allocation across observationally identical individuals. Under this interpretation, the capacity constraint can still be satisfied ex-post.

Sun (2024) studies a binary treatment rule $g(X) \in \{0, 1\}$ under an unknown budget constraint. Because g is binary, the allocation itself is deterministic. However, the rule may still violate the (true) budget constraint because the constraint is not known exactly. Although our framework shares with hers the possibility of ex-post deviations from the intended constraint, the source of the problem is different, and the two frameworks are not the same.

Remark 2.5. Although this paper focuses on capacity constraints, the OT formulation can accommodate additional constraints beyond fixing the treatment marginal F_T . Any finitely many linear constraints on the coupling μ of the form $\int f_i d\mu \leq c_i$ (or $= c_i$) with bounded continuous f_i preserve the compact convexity of the feasible set,⁵ so the theoretical framework developed in Section 3 extends to such settings. Two examples of practical interest are:

- (1) *Demographic parity.* Suppose the covariate $X = (A, X_1)$ includes a sensitive attribute $A \in \{a, b\}$ (e.g., gender) and other characteristics X_1 . The constraint

$$|P(T = 1|A = a) - P(T = 1|A = b)| \leq \varepsilon$$

bounds the disparity in treatment rates across groups. Since $P(T = 1|A = g) = \mu(\mathcal{T}_1 \times \{A = g\} \times \mathcal{X}_1) / P(A = g)$ for $g = a, b$ and $P(A = g)$ is determined by F_X , this constraint is linear in μ and can be incorporated directly.

- (2) *Covariate-blind fairness.* Requiring that treatment assignment does not depend on the sensitive attribute conditional on other covariates—i.e., $\mu(t|a, x_1) = \mu(t|b, x_1)$ for all t and x_1 —is also a linear constraint on μ , since F_X is fixed.

In principle, the theoretical analysis in Section 3 extends directly to such settings by replacing \mathcal{M} with the constrained coupling space $\mathcal{M}' \subset \mathcal{M}$. The optimal rules over \mathcal{M}' are constrained-optimal and may differ from the first-best rules over \mathcal{M} . However, the practical implementation requires solving an optimal transport problem with additional linear constraints, for which standard OT solvers (designed for marginal constraints only) do not directly apply. Developing efficient computational methods for such constrained OT problems in the treatment choice context is left for future work.

2.3. The data generating process. When the true (finite-dimensional) parameter θ_0 is known to the planner, the optimal rules can be obtained by solving

$$\max_{\mu \in \mathcal{M}} W(\theta_0, \mu).$$

However, since the planner does not know θ_0 in practice, she must select a rule in a data-driven manner. For this purpose, data $Z^n = (Z_1, \dots, Z_n)$, which are informative about θ (and hence about $F_t(\cdot|x, \theta)$) are available from a training population. We assume that the data Z^n are i.i.d. with $Z_i \sim P_\theta$ on some space \mathcal{Z} equipped with the Borel σ -algebra $\mathcal{B}(\mathcal{Z})$. We let P_θ^n denote the joint probability measure of Z^n . In Appendix G, we consider an extension in which the sampling distribution of data may depend on (possibly infinite-dimensional) nuisance parameters, as in a GMM model.

⁵Convexity is preserved because the constraint sets are convex (as sublevel sets of linear functionals) and the intersection of convex sets is convex. Compactness is preserved because $\mu \mapsto \int f_i d\mu$ is continuous in the Wasserstein topology when f_i is bounded continuous (since $\mathcal{X} \times \mathcal{T}$ is compact, the Wasserstein topology coincides with weak convergence), so each constraint set is closed, and a closed subset of the compact set \mathcal{M} is compact.

Example 2.6. Angrist *et al.* (2006) study the medium-term effects of the PACES program in Columbia.⁶ Specifically, they investigate the impact of receiving a randomly assigned voucher (which allows attendance at private schools) on educational attainment seven years later, using the administrative data.

In one of their main specifications, they consider the following linear model:

$$Y_i = X_i^\top \beta + \alpha T_i + u_i,$$

where Y_i denotes test scores, X_i includes observable covariates (sex and age), T_i is an indicator for the treatment status, and u_i is an error term. However, not all individuals in the sample took the exam. Because the voucher recipients were more likely to take the exam than non-recipients, a selection issue arises. To address this, Angrist *et al.* (2006) construct a modified test score variable by censoring the observed scores at a specific quantile of the test score distribution. Let R_i be an indicator for exam registration and $\tau > 0$ be the censoring threshold. Then the censored test score is defined as

$$Y_i(\tau) := \mathbf{1}\{R_i Y_i \geq \tau\} Y_i + \mathbf{1}\{R_i Y_i < \tau\} \tau.$$

Under the assumptions that (i) u_i is normally distributed with mean zero, and (ii) any untested student would have scored below the threshold τ had they taken the exam, the parameters can be consistently estimated using a Tobit-type maximum likelihood estimator. In this context, the parameter is $\theta = (\alpha, \beta, \sigma)$, and the observed data is $Z^n = \{(T_i, X_i, Y_i(\tau)) : i = 1, \dots, n\}$, where σ is the standard deviation of the error term u_i . ■

Following Hirano and Porter (2009) and among others, we use a local asymptotic framework where we perturb the data-generating process around the true one. Let Θ be an open subset of \mathbb{R}^k and suppose that θ_0 is the true parameter. Let $\theta_{nh} := \theta_0 + h/\sqrt{n}$. We assume that the sequence of experiments $\mathcal{E}_n = \{P_\theta^n : \theta \in \Theta\}$ satisfies *differentiability in quadratic mean* (DQM) at θ_0 : there exists a function $s : \mathcal{Z}^n \rightarrow \mathbb{R}^k$ such that

$$\int \left[dP_{\theta_0+h}^{1/2}(z) - dP_{\theta_0}^{1/2}(z) - \frac{1}{2} h' s(z) dP_{\theta_0}^{1/2}(z) \right]^2 = o(\|h\|^2) \quad \text{as } h \rightarrow 0, \quad (2.5)$$

where s is the score function associated with \mathcal{E}_1 . Let $I_0 = \mathbb{E}_{P_{\theta_0}^n} [ss']$.

The planner's *statistical treatment assignment rule* (or just *rule*) $\mu : \mathcal{Z}^n \rightarrow \mathcal{M}$ maps realizations of data into the coupling. Let

$$A_0 := \arg \max_{\mu \in \mathcal{M}} W(\theta_0, \mu)$$

be the set of couplings that maximize the welfare at θ_0 . We define the class of sequences of rules by

$$\mathcal{D} := \left\{ \{\mu_n\} : \mu_n(Z^n) \overset{h}{\rightsquigarrow} Q_{\theta_0, h} \text{ and } \sqrt{n} P_{\theta_{nh}}^n \left(\mu_n(Z^n) \notin A_0^\delta \right) \rightarrow 0 \quad \forall h \in \mathbb{R}^k, \forall \delta > 0, \forall \theta_0 \in \Theta \right\}, \quad (2.6)$$

⁶PACES stands for Programa de Ampliación de Cobertura de la Educación Secundaria.

where $\overset{h}{\rightsquigarrow}$ denotes convergence in distribution along $P_{\theta_{nh}}^n$ with $Z^n \sim P_{\theta_{nh}}^n$ for each n , and $Q_{\theta_0, h}$ is a (possibly degenerate) probability measure on \mathcal{M} . For technical reasons, we require $\sqrt{n}P_{\theta_{nh}}^n(\mu_n(Z^n) \notin A_0^\delta) \rightarrow 0$ for all $\delta > 0$. This condition ensures that the treatment rule maximizes the welfare at the true parameter θ_0 with high probability in a neighborhood of θ_0 , and that the probability of selecting a suboptimal coupling (i.e., $\mu_n \notin A_0$) vanishes sufficiently fast.⁷

Since (\mathcal{M}, d_W) is compact (and thus complete and separable) by Villani (2009, Remark 6.19), we obtain the following asymptotic representation theorem by van der Vaart (1991, Theorem 3.1).

Proposition 2.7. *Let $\{\mu_n\} \in \mathcal{D}$ satisfy $\mu_n \overset{h}{\rightsquigarrow} Q_{\theta_0, h}$ for all $h \in \mathbb{R}^k$ and $\theta_0 \in \Theta$. Under Assumption 3.1 (i)–(iii) below, there exists $\mu_\infty : \mathbb{R}^k \times [0, 1] \rightarrow \mathcal{M}$ such that for every $h \in \mathbb{R}^k$, $\mathcal{L}_h(\mu_\infty(\Delta, U)) = Q_{\theta_0, h}$, $\mathcal{L}_h(\Delta) = N(h, I_0^{-1})$, and $\mathcal{L}_h(U) = \text{Unif}[0, 1]$ with $U \perp\!\!\!\perp \Delta$.*

This proposition states that any sequence $\{\mu_n\}$ in \mathcal{D} can be matched by some treatment rule μ_∞ in a *limit experiment* where we observe a single draw Δ from a shifted normal distribution and an independent uniform random variable U . This representation is useful for analyzing the asymptotic properties of rules, as the limit experiment is more analytically tractable than the original experiments \mathcal{E}_n .

Remark 2.8. Hirano and Porter (2009) study the optimality of treatment assignment rules after conditioning on a fixed covariate value $X = x$. In our notation, this is equivalent to finding an optimal conditional probability $\mu(\cdot|x)$ for each x . Accordingly, they derive the asymptotic representation of $\mu(\cdot|x)$ as a function of Z^n by applying a version of the representation theorem (see their Proposition 3.1). However, because it is difficult to accommodate the capacity constraints within this framework, we instead apply the representation theorem to couplings $\mu \in \mathcal{M}$. Note that in our setup, the map $Z^n \mapsto \mu(Z^n)$ takes values in \mathcal{M} , rather than the unit interval.

3. OPTIMAL DECISIONS

3.1. Decision theoretic framework and rules. We begin by introducing a decision theoretic framework to evaluate the performance of a sequence of rules $\{\mu_n\} \in \mathcal{D}$. For each parameter value θ , let

$$W_{\mathcal{M}}^*(\theta) := \max_{\mu' \in \mathcal{M}} W(\theta, \mu')$$

denote the maximum attainable welfare, and let $\dot{w}_{\theta_0}(x, t; h)$ denote the directional derivative of $w(\theta, x, t)$ at θ_0 in direction h .

In the literature on locally asymptotically optimal treatment rules, decision rules are typically evaluated by integrating the local risk over perturbations of the parameter. One example of such a criterion is

⁷This condition is weaker than $\sqrt{n}P_{\theta_{nh}}^n(\mu_n(Z^n) \notin A_0) \rightarrow 0$. Since $\bigcap_{\delta > 0} A_0^\delta = A_0$, this still ensures that the asymptotic distribution of any rule in \mathcal{D} has support in A_0 . When \mathcal{M} is finite, as in Christensen *et al.* (2026), the two conditions are equivalent because A_0 is separated from its complement.

$$\inf_{\{\mu_n\} \in \mathcal{D}} \int_{\mathbb{R}^k} \liminf_{n \rightarrow \infty} \sqrt{n} R(\mu_n, \theta_{nh}) dh, \quad (3.1)$$

where

$$R(\mu, \theta) := \mathbb{E}_{P_\theta^n} [W_{\mathcal{M}}^*(\theta) - W(\theta, \mu(Z^n))] \quad (3.2)$$

is the welfare-regret risk associated with the decision rule $Z^n \mapsto \mu(Z^n) \in \mathcal{M}$ at θ . This average-optimality criterion is similar in spirit to the one employed by Christensen *et al.* (2026), whereas Hirano and Porter (2009) and Xu (2024) consider a related formulation in which the integration over h is weighted by a prior density.

A major difficulty with the classical average-risk criterion is that

$$\int_{\mathbb{R}^k} \liminf_{n \rightarrow \infty} \sqrt{n} R(\mu_n, \theta_{nh}) dh$$

may be infinite for every sequence $\{\mu_n\} \in \mathcal{D}$. This can occur because $\dot{w}_{\theta_0}(x, t; h)$ may grow too quickly for the integrand to be integrable.

To obtain an optimality notion that is always well-defined, we turn to the Gaussian limit experiment. By Proposition 2.7, any $\{\mu_n(Z^n)\} \in \mathcal{D}$ is matched by some coupling $\mu_\infty(\Delta, U)$ in the limit experiment. Motivated by this representation, we evaluate rules through their induced decisions in the Gaussian shift experiment.

Definition 3.1. A sequence $\{\mu_n\} \in \mathcal{D}$ is said to be *optimal in the limit experiment* if, for each fixed $h \in \mathbb{R}^k$, it converges in distribution along $P_{\theta_{nh}}^n$ to a coupling μ_∞ that solves

$$\min_{\mu \in A_0} L_\infty(\mu, \Delta) \quad (3.3)$$

for any realization of signal $\Delta \sim N(h, I_0^{-1})$, where the limiting loss is defined by

$$L_\infty(\mu, \Delta) := \int \left[\dot{W}_{\mathcal{M},0}^*[h] - \int \dot{w}_{\theta_0}(x, t; h) d\mu \right] dN(\Delta, I_0^{-1})(h),$$

and $\dot{W}_{\mathcal{M},0}^*(h) = \max_{\mu \in \mathcal{M}} \int \dot{w}_{\theta_0}(x, t; h) d\mu$ denotes the Hadamard directional derivative of $W_{\mathcal{M}}^*$ at θ_0 .⁸

The distribution $N(\Delta, I_0^{-1})$ corresponds to the posterior distribution of the local parameter h in the Gaussian shift experiment under a flat prior. Thus $L_\infty(\mu, \Delta)$ is interpreted as the limiting posterior risk of choosing coupling μ upon observing signal Δ .

We next briefly describe the relationship between the classical average-risk criterion and optimality in the limit experiment; a formal statement is deferred to Section 3.5.

One can show (Theorem 3.10) that, under additional regularity conditions, a sequence $\{\mu_n^*\} \in \mathcal{D}$ is average optimal if $\{\mu_n^*\}$ is matched by some coupling μ_∞^* in the limit experiment such that μ_∞^* solves

$$\min_{\mu \in A_0} \int \int \left[\dot{W}_{\mathcal{M},0}^*[h] - \int \dot{w}_{\theta_0}(x, t; h) d\mu \right] dN(h, I_0^{-1})(\Delta) dh.$$

⁸The existence of the derivative $\dot{W}_{\mathcal{M},0}^*(h)$ is formally shown in Lemma B.1.

Because the integrand is nonnegative, we may exchange the order of integration. Doing so yields an optimization problem that coincides exactly with the limit-experiment optimality condition (3.3).

This connection is standard in local asymptotic decision theory: Hirano and Porter (2009) use an argument that minimizing posterior risk in the Gaussian limit experiment yields average optimality, and the subsequent literature adopts arguments that are, in essence, solving the corresponding limit-experiment problem. Our analysis demonstrates that this classical connection continues to hold even in our non-standard action space—the space of couplings (\mathcal{M}, d_W) arising from capacity constraints and the optimal-transport formulation of the planner’s problem.

In contrast to the classical average-optimality criterion, the limit-experiment optimality criterion is always well-defined, and captures the condition that any average-optimal rule would have to satisfy *if* the classical criterion were finite. Thus, even when the average risk is not well-defined, optimality in the limit experiment provides a sharp and meaningful approximation to the notion of average optimality.

3.2. Assumptions. As in Clarke and Barron (1990) and Christensen *et al.* (2026), we say that a family \mathcal{P} is *locally quadratic* if for any $\theta_0 \in \Theta$, $D_{\text{KL}}(p_\theta \parallel p_{\theta'}) \leq 2(\theta - \theta')^\top I_0(\theta - \theta')$ holds for any θ and θ' belonging to a neighborhood of θ_0 , where $D_{\text{KL}}(p_\theta \parallel p_{\theta'})$ is the Kullback-Leibler divergence with respect to a common dominating measure ν . Also, we say \mathcal{P} is *sound* if weak convergence of P_θ to $P_{\theta'}$ is equivalent to the convergence of θ to θ' for probability measures $P_\theta, P_{\theta'}$ and parameters $\theta, \theta' \in \Theta$.

Assumption 3.1. (i) Θ is open. (ii) \mathcal{P} is DQM at any θ_0 . (iii) I_0 is finite and nonsingular at any θ_0 . (iv) \mathcal{P} is locally quadratic. (v) \mathcal{P} is sound.

The first three conditions are standard assumptions in local asymptotic frameworks. Conditions (iv) and (v) ensure that Schwartz’s theorem—which originally establishes posterior consistency in a space of density functions (see, for example, Ghosh and Ramamoorthi (2003) and Ghosal and van der Vaart (2017))—can be applied in a parametrized setting, as in Clarke and Barron (1990). We impose these assumptions to show that the Bayesian rule $\{\mu_n^B\}$ satisfies $\sqrt{n}P_{\theta_{nh}}^n(\mu_n^B(Z^n) \notin A_\delta) \rightarrow 0$ for all $\delta > 0$. The same conditions are also imposed by Christensen *et al.* (2026).

Assumption 3.2. $\mathcal{X} \times \mathcal{T}$ is compact in the product metric space where the distance function d is equipped.

For example, this condition is satisfied if \mathcal{X} is a compact metric space, and \mathcal{T} is a finite discrete space.

Assumption 3.3. (i) $w(\theta, x, t)$ is bounded continuous on Θ for any $(x, t) \in \mathcal{X} \times \mathcal{T}$. (ii) $w(\theta, x, t)$ is continuous on $\mathcal{X} \times \mathcal{T}$ uniformly over Θ .

Note that discrete covariates are compatible with condition (ii), since the metric d can be defined to incorporate the discrete metric.

Assumption 3.4. (i) $w(\theta, \cdot)$ is directionally differentiable (as a function on $\mathcal{X} \times \mathcal{T}$) at any $\theta \in \Theta$ with derivative \dot{w}_θ .⁹ (ii) $\dot{w}_{\theta_0}(x, t; h)$ is continuous on $\mathcal{X} \times \mathcal{T}$ for any h . (iii) $\dot{w}_{\theta_0}(\cdot; h)$ is uniformly dominated by a function $K(h)$ that grows at most subpolynomially of order p ; i.e., $\max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\dot{w}_{\theta_0}(x, t; h)| \leq K(h) \leq 1 + \|h\|^p$.

Condition (i) imposes a uniform version of directional differentiability, rather than requiring it only pointwise in (x, t) . For condition (iii), a similar polynomial growth condition appears in the study of Bayes estimators (van der Vaart, 1998, Section 10.3). Choosing $p = 1$ is sufficient for (2.1) provided $\max_{(x,t) \in \mathcal{X} \times \mathcal{T}} \left\| \frac{\partial}{\partial \theta} \int y dF_t(y|x, \theta_0) \right\| < \infty$. Condition (ii) is used in Theorem 3.4.

Assumption 3.5. (i) The prior (Lebesgue) density function π is positive, continuous, and bounded on Θ . (ii) $\int \|\theta\|^p d\pi(\theta) < \infty$.

The order p used in condition (ii) must align with the order in Assumption 3.4 (iii). Condition (ii) is used in Theorem 3.10.

In the following subsections, we investigate two canonical assignment rules to check whether they are optimal in the limit experiment. The first rule is the *plug-in rule* $\{\mu_n^P(Z^n)\}$, defined for each n by

$$\mu_n^P(Z^n) \in \arg \max_{\mu \in \mathcal{M}} W(\hat{\theta}_n, \mu),$$

where $\hat{\theta}_n$ is a *best regular estimator* of θ_0 such that

$$\sqrt{n}(\hat{\theta}_n - \theta_{nh}) \overset{h}{\rightsquigarrow} N(0, I_0^{-1}) \quad \text{as } n \rightarrow \infty.$$

The maximum likelihood estimator or the Bayesian posterior mean estimator are typical examples of best regular estimators. From the arguments of Hirano and Porter (2009), such a plug-in rule is shown to be optimal in the limit experiment in binary treatments setup without constraints on available treatments.

The second rule is the *Bayesian rule* $\{\mu_n^B(Z^n)\}$, defined for each n by

$$\mu_n^B(Z^n) \in \arg \max_{\mu \in \mathcal{M}} \int_{\Theta} W(\theta, \mu) \pi_n(\theta | Z^n) d\theta,$$

where $\pi_n(\theta | Z^n)$ is the posterior density obtained from a strictly positive, continuous prior density π on Θ . From the arguments of Christensen *et al.* (2026), such a Bayesian rule is shown to be optimal in the limit experiments in discrete choice problems when (i) the model is partially identified, and (ii) decision rules are not fractional.

3.3. Asymptotic behavior of the Bayesian rule. It should be noted that the Bayesian rule $\{\mu_n^B\}$ may not be uniquely determined as our framework allows for multiple maximizers of the objective

⁹That is, for any $r_n \downarrow 0$ and $h_n \rightarrow h$,

$$\max_{(x,t) \in \mathcal{X} \times \mathcal{T}} \left| \frac{w(\theta_0 + r_n h_n, x, t) - w(\theta_0, x, t)}{r_n} - \dot{w}_{\theta_0}(x, t; h) \right| \rightarrow 0.$$

function. This non-uniqueness complicates the analysis since we cannot directly apply the argmax theorem, which is often used to study the asymptotic behavior of general argmax-functionals.

To deal with this, we utilize a penalized version of the Bayesian rule. Let $\nu \in \mathcal{M}$ be any fixed reference measure and $\lambda : \mathcal{M} \rightarrow \mathbb{R}_+$ be a functional given by $\mu \mapsto (d_W(\mu, \nu))^2$, which will serve as a penalty function of a maximization problem. For example, we can let ν be the product measure of F_X and F_T . The functional λ has following properties:

Proposition 3.2. *H is a nonnegative, continuous, strictly convex, and bounded functional on (\mathcal{M}, d_W) .*

Proof. See Appendix E. □

Then we define the penalized Bayesian rule by

$$\mu_{n,\varepsilon}^B(z) := \arg \max_{\mu \in \mathcal{M}} \int \sqrt{n}W(\theta, \mu)\pi_n(\theta|z)d\theta - \varepsilon H(\mu),$$

for $\varepsilon > 0$. Note that $\mu_{n,\varepsilon}^B$ becomes the unique maximizer of this penalized problem by the strict convexity of H . We then obtain a useful result on the penalized rules $\{\mu_{n,\varepsilon}^B\}$ by following the arguments of Nutz (2022).¹⁰

Proposition 3.3. *Let $\mathcal{M}_{opt}(z) := \arg \max_{\mu \in \mathcal{M}} \int \sqrt{n}W(\theta, \mu)\pi_n(\theta|z)d\theta$. For each $z \in \mathcal{Z}^n$, there exists a unique $\mu_n^B(z) \in \mathcal{M}$ such that (i) $\mu_{n,\varepsilon}^B(z)$ converges weakly to $\mu_n^B(z)$ as $\varepsilon \downarrow 0$, (ii) $\mu_n^B(z) \in \mathcal{M}_{opt}(z)$, and (iii) $\mu_n^B(z) = \arg \min_{\mu \in \mathcal{M}_{opt}(z)} H(\mu)$.*

Proof. See Appendix F. □

Thus, we can construct a unique sequence $\{\mu_n^B\}$ by defining $\mu_n^B(z)$ as the minimizer of the penalty function H over $\mathcal{M}_{opt}(z)$. All results on the Bayesian rule below are stated for the sequence $\{\mu_n^B\}$ constructed in this way.

Theorem 3.4. *Under Assumptions 3.1–3.5, $\{\mu_n^B\}$ is optimal in the limit experiment.*

Proof. See Appendix A. □

Remark 3.5. Formally, the construction of $\{\mu_n^B\}$ depends on the choice of a prior distribution. However, any prior satisfying Assumption 3.5 leads to the same conclusion.

Remark 3.6. $\{\mu_n^B\}$ remains optimal in the limit experiment if the directional differentiability in Assumption 3.4 is strengthened to the full differentiability.

¹⁰Nutz (2022) provides corresponding results by choosing H as the Kullback-Leibler (KL) information criterion between $\mu \in \mathcal{M}$ and any reference measure in \mathcal{M} . KL is nonnegative and strictly convex in μ , but not continuous and bounded. Here we impose stronger requirements for the penalty function H , which is needed to handle weak convergence of functionals on \mathcal{M} to study the asymptotic properties of rules. Accordingly, the mode of convergence of $\mu_{n,\varepsilon}^B(z)$ is modified to weak convergence from convergence in total variation, see Nutz (2022, Theorem 5.5).

Remark 3.7. Christensen *et al.* (2026) impose a condition called *no first-order ties*, which requires the uniqueness of the minimizer of the loss function in the limit experiment. This condition addresses an indeterminacy: their treatment rule must be matched to one of the minimizers, but it is not uniquely determined without this condition. By contrast, our construction of the rule $\{\mu_n^B(Z^n)\}$ allows us to avoid imposing this condition, since $\mu_n^B(z)$ is matched with μ_∞^* in the limit experiment where μ_∞^* uniquely maximizes the penalty function H over the set of minimizers of $L_\infty(\mu, \Delta)$.¹¹

3.4. Asymptotic behavior of the plug-in rule.

3.4.1. *When w is directionally differentiable.* From the definition of the limit-experiment optimality, it is easy to see that any matched rule must solve

$$\max_{\mu \in A_0} \int \left[\int \dot{w}_{\theta_0}(x, t; h) d\mu \right] dN(\Delta, I_0^{-1})(h). \quad (3.4)$$

One can show that $\{\mu_n^P\} \in \mathcal{D}$, which implies the asymptotic distribution of $\mu_n^P(Z^n)$ along $P_{\theta_{nh}}^n$ has a support in A_0 . Thus, for sufficiently large n , $\mu_n^P(Z^n)$ equivalently solves

$$\begin{aligned} \arg \max_{\mu \in A_0} \sqrt{n} W(\hat{\theta}_n, \mu) &= \arg \max_{\mu \in A_0} \int \sqrt{n} w(\hat{\theta}_n, x, t) d\mu \\ &= \arg \max_{\mu \in A_0} \sqrt{n} \left[\int w(\hat{\theta}_n, x, t) d\mu - \int w(\theta_0, x, t) d\mu \right] \end{aligned}$$

where the second equality follows because the value of $W(\theta_0, \mu)$ is constant across $\mu \in A_0$. We will see that the maximization problem that plug-in rules solve weakly converges to a different maximization problem from (3.4). We actually claim that

$$\max_{\mu \in A_0} \sqrt{n} \left[\int w(\hat{\theta}_n, x, t) d\mu - \int w(\theta_0, x, t) d\mu \right] \xrightarrow{h} \max_{\mu \in A_0} \int \dot{w}_{\theta_0}(x, t; \Delta) d\mu \quad \text{as } n \rightarrow \infty, \quad (3.5)$$

where $\Delta \sim N(h, I_0^{-1})$.

To see this, let $B_n(\mu) := \sqrt{n} \left[\int w(\hat{\theta}_n, x, t) d\mu - \int w(\theta_0, x, t) d\mu \right]$ and $B_\infty(\mu) := \int \dot{w}_{\theta_0}(x, t; \Delta) d\mu$. To simplify the argument, we impose a high-level condition that the process $\{B_n(\mu) : \mu \in A_0\}$ is asymptotically tight.

By the best regularity of $\hat{\theta}_n$, it follows that $\sqrt{n}(\hat{\theta} - \theta_0) = I_0^{-1} S_n + o_{P_{\theta_0}^n}(1)$ with $S_n \xrightarrow{0} N(0, I_0)$ as $n \rightarrow \infty$. Combining Le Cam's third lemma and the delta method for the directionally differentiable functions (Fang and Santos, 2019, Theorem 2.1) yields

$$\sqrt{n} \left[\int w(\hat{\theta}_n, x, t) d\mu - \int w(\theta_0, x, t) d\mu \right] \xrightarrow{h} \int \dot{w}_{\theta_0}(x, t; \Delta) d\mu \quad \text{as } n \rightarrow \infty,$$

where $\Delta \sim N(h, I_0^{-1})$. By the asymptotic tightness of $\{B_n(\mu) : \mu \in A_0\}$, we can extend this result to convergence in distribution of the process

$$B_n \xrightarrow{h} B_\infty \quad \text{as } n \rightarrow \infty \text{ on } \ell^\infty(A_0),$$

¹¹Xu (2024) assumes uniqueness of the rule and therefore does not encounter the issue of non-uniqueness we addressed here.

by van der Vaart and Wellner (1996, Theorem 1.5.4). Then applying the continuous mapping theorem yields

$$\max_{\mu \in A_0} B_n(\mu) \xrightarrow{h} \max_{\mu \in A_0} B_\infty(\mu) \quad \text{as } n \rightarrow \infty,$$

which completes the argument.

It is evident that the solutions of RHS of (3.5) need not to solve (3.4). Thus the plug-in rules might not be optimal in the limit experiment in general when w is directionally differentiable. In the following, we provide a concrete example such that the plug-in rule becomes sub-optimal.

Example 3.8. Assume that the planner observes a single covariate X , say sex, with $\mathcal{X} = \{x_f, x_m\}$. Suppose $F_X(x_f) = F_X(x_m) = 1/2$. Consider binary treatments setup where a fraction $p = 1/2$ of the individuals to be treated. Assume that the conditional mean of the potential outcome is given by

$$\int y dF_t(y|x, \theta) = \begin{cases} t\theta^2 & \text{if } x = x_f, \\ t\theta & \text{if } x = x_m, \end{cases} \quad \theta_0 = 1.$$

Thus, the mean outcome is identical for x_f and x_m at true value θ_0 , but differs locally around θ_0 . Further, let $\varepsilon > 0$ satisfy $-\varepsilon < y_\ell(0)$. In this case, the directional derivative of the maximin utility function defined in (2.1) with $\lambda = 0$ is

$$\dot{w}_{\theta_0}(x_f, 1; h) = \begin{cases} 2h & \text{if } h \geq 0, \\ 0 & \text{if } h < 0, \end{cases} \quad \dot{w}_{\theta_0}(x_f, 0; h) = 0,$$

and

$$\dot{w}_{\theta_0}(x_m, 1; h) = \begin{cases} h & \text{if } h \geq 0, \\ 0 & \text{if } h < 0, \end{cases} \quad \dot{w}_{\theta_0}(x_m, 0; h) = 0.$$

Note that $\mathcal{M} = A_0$ since $\theta_0 = 1$. For simplicity, assume $I_0 = 1$.

Under this setup, we have

$$\int \int \dot{w}_{\theta_0}(x, t; s) dN(\Delta, I_0^{-1})(s) d\mu = \frac{1}{2} h(\Delta) \mu(1|x_f) + h(\Delta) p, \quad h(\Delta) := \left[\Delta - \frac{\phi(-\Delta)}{1 - \Phi(-\Delta)} \right],$$

where ϕ and Φ denote the pdf and cdf of the standard normal distribution, respectively. Let Δ_h satisfy $h(\Delta_h) = 0$ (numerically, $\Delta_h \approx 0.506$). Then note that $h(\Delta) \geq 0$ if $\Delta \geq \Delta_h$ and $h(\Delta) < 0$ if $\Delta < \Delta_h$. Therefore, the optimal rule in (3.4) is

$$\begin{cases} \mu(1|x_f) = 1, \mu(1|x_m) = 0 & \text{if } \Delta \geq \Delta_h, \\ \mu(1|x_f) = 0, \mu(1|x_m) = 1 & \text{if } \Delta < \Delta_h. \end{cases}$$

In contrast, we obtain

$$\int \dot{w}_{\theta_0}(x, t; \Delta) d\mu = \left[\max\{\Delta, 0\} - \max\left\{\frac{\Delta}{2}, 0\right\} \right] \mu(1|x_f) + \max\{\Delta, 0\} p.$$

Therefore, the optimal rule in (3.5) is

$$\begin{cases} \mu(1|x_f) = 1, \mu(1|x_m) = 0 & \text{if } \Delta \geq 0, \\ \mu(1|x_f) = 0, \mu(1|x_m) = 1 & \text{if } \Delta < 0, \end{cases}$$

which is suboptimal for $0 \leq \Delta \leq \Delta_h$. ■

Remark 3.9. To investigate the asymptotic properties of the plug-in rule formally, we need to handle the non-uniqueness issue. This can be done by the penalization used for the construction of the Bayesian rule.

3.4.2. *When w is fully differentiable.* If we strengthen the directional differentiability in Assumption 3.4 (i) to the continuous differentiability, then the plug-in rules attain the optimal value (3.4). To see this, notice that the continuous differentiability of $w(\theta_0, x, t)$ implies that the directional derivative is linear; i.e., $\dot{w}_{\theta_0}(x, t; s) = \dot{w}_{\theta_0}(x, t)^\top s$ for some $\dot{w}_{\theta_0}(x, t) \in \mathbb{R}^k$. Then (3.4) becomes

$$\max_{\mu \in A_0} \int \dot{w}_{\theta_0}(x, t)^\top \Delta d\mu,$$

which is the same as (3.5)

This pattern is consistent with findings from the existing literature. As discussed earlier, the plug-in rule is optimal in point-identified models when the utility function is continuously differentiable. In contrast, Christensen *et al.* (2026) demonstrate that, in partially identified models, the Bayesian rule is optimal when the utility function is only directionally differentiable, whereas the plug-in rule fails to attain the optimal value for each signal Δ in the limit experiment unless full differentiability holds. In partially identified settings, directional differentiability is a natural and often unavoidable assumption, as full differentiability typically does not hold.

3.5. Connection to average optimality. We now formalize the connection between limit-experiment optimality and the classical average-risk criterion. The following result shows that, under an additional condition, limit-experiment optimality implies average optimality.

Theorem 3.10. *Let $\{\mu_n\} \in \mathcal{D}$ be any sequence of decision rules that is matched by μ_∞ in the limit experiment, and suppose additionally that $\sqrt{n}P_{\theta_{nh}}^n(\mu_n(Z^n) \notin A_0) \rightarrow 0$ for all h and θ_0 . Under Assumptions 3.1–3.5, $\{\mu_n\} \in \mathcal{D}$ is average optimal if, for each fixed $h \in \mathbb{R}^k$, $\mu_\infty(\Delta) \in \arg \min_{\mu \in A_0} L_\infty(\mu, \Delta)$ for any realization of the signal $\Delta \sim N(h, I_0^{-1})$, whenever the average optimality is well-defined.*

Proof. See Appendix C. □

Remark 3.11. The additional condition $\sqrt{n}P_{\theta_{nh}}^n(\mu_n \notin A_0) \rightarrow 0$ is stronger than the condition $\sqrt{n}P_{\theta_{nh}}^n(\mu_n \notin A_0^\delta) \rightarrow 0$ for all $\delta > 0$ that defines \mathcal{D} . The following corollary provides sufficient conditions under which the Bayesian rule $\{\mu_n^B\}$ is average optimal without imposing this additional condition.

For the next result, we strengthen Assumption 3.3 as follows.

Assumption 3.6. (i) $w(\theta, x, t)$ is Lipschitz continuous on $\mathcal{X} \times \mathcal{T}$ uniformly over Θ with Lipschitz constant L ; i.e., $|w(\theta, x, t) - w(\theta, x', t')| \leq Ld((x, t), (x', t'))$ for all $\theta \in \Theta$ and $(x, t), (x', t') \in \mathcal{X} \times \mathcal{T}$. (ii) $w(\theta, x, t)$ is Hölder continuous in θ of order $\alpha \in (0, 1]$ uniformly over $\mathcal{X} \times \mathcal{T}$; i.e., there exists $C > 0$ such that $|w(\theta, x, t) - w(\theta', x, t)| \leq C \|\theta - \theta'\|^\alpha$ for all $(x, t) \in \mathcal{X} \times \mathcal{T}$ and $\theta, \theta' \in \Theta$.

Note that Assumption 3.6 (i) implies Assumption 3.3 (ii), and Assumption 3.6 (ii) implies Assumption 3.3 (i). When w is continuously differentiable in both θ and (x, t) on the compact domain $\Theta \times \mathcal{X} \times \mathcal{T}$, Assumption 3.6 holds with $\alpha = 1$.

Corollary 3.12. Under Assumptions 3.1–3.5 and Assumption 3.6, suppose that the gap function $\eta(\delta) := \min_{\mu \in A_0} W(\theta_0, \mu) - \sup_{\nu \notin A_0^\delta} W(\theta_0, \nu)$ satisfies $\eta(\delta) \geq c\delta^\gamma$ for some $c > 0$ and $\gamma \in (0, 2\alpha)$ for all sufficiently small $\delta > 0$. Then the Bayesian rule $\{\mu_n^B\}$ is average optimal whenever the average optimality is well-defined.

Proof. See Appendix D. □

4. SIMULATION

We conduct a simulation study to evaluate the performance of the Bayesian rule and the plug-in rule under the following conditions: (i) the welfare function is either smooth or only directionally differentiable, and (ii) the sample size is relatively small ($n = 200$) and large ($n = 500$).

We closely follows the data generating process described in Example 2.6. For the training population, the latent variable is generated by

$$Y_i^* = X_i^\top \beta + \alpha T_i + u_i,$$


where $X_i \in \mathbb{R}^2$ denotes the observable covariates and T_i is the binary treatment that is randomly assigned. The first coordinate of X_i , interpreted as age, follows a truncated normal distribution with mean 4, standard deviation of 2, and is bounded on $[1, 10]$. The second coordinate, interpreted as sex, is a binary variable assigned with equal probability. The observed outcome is

$$Y_i = \max\{0, Y_i^*\}.$$

We set $\beta_0 = (-2, -3)$, $\alpha_0 = 4$, and $u_i \sim N(0, \sigma_0^2)$ with $\sigma_0 = 10$. The observed data is an i.i.d. sample $Z^n = \{(Y_i, X_i, T_i)\}_{i=1}^n$. The parameters $\theta_0 = (\beta_0, \alpha_0, \sigma_0)$ can be estimated by the maximum likelihood using Z^n .

In this Tobit model, the conditional mean of the potential outcomes in the training population is given by

$$w(\theta, x, t) = (x^\top \beta + \alpha t) - (x^\top \beta + \alpha t) \Phi \left(\frac{-x^\top \beta - \alpha t}{\sigma} \right) + \sigma \phi \left(\frac{-x^\top \beta - \alpha t}{\sigma} \right),$$

FIGURE 1. Welfare contrasts under smooth and directionally differentiable welfare at θ_0 

where Φ and ϕ denote the standard normal cdf and pdf, respectively. Following Examples 2.1 and 2.2, we specify the planner's utility function as

$$w_R(\theta, x, t, \varepsilon, \lambda) = \lambda w(\theta, x, t) + (1 - \lambda) \max \{w(\theta, x, t) - \varepsilon, 0\},$$

where $\lambda \in [0, 1]$ and $\varepsilon > 0$. We note that w_R is differentiable when $\lambda = 1$, but only directionally differentiable otherwise. In what follows, we focus on the cases $\lambda = 0, 1$ and $\varepsilon = 0.8$.

Figure 1 plots the welfare contrasts $w_R(\theta, x, 1, \varepsilon, \lambda) - w_R(\theta, x, 0, \varepsilon, \lambda)$ for both males and females at θ_0 . Under $\lambda = 0$, kinks appear when $w(\theta, x, t) - \varepsilon < 0$. For a given covariate x , the contrast is zero if both $w(\theta, x, 1) < \varepsilon$ and $w(\theta, x, 0) < \varepsilon$; that is, individuals with sufficiently low welfare are regarded as deriving no benefit from treatment.

The upper panels of Figure 4 show the oracle (infeasible optimal) rule at θ_0 . The rule assigns treatment to females younger than approximately 6.5 and to males younger than approximately 5. Notably, this oracle rule remains unchanged across $\lambda = 0$ and $\lambda = 1$. It also remains the same at $\theta_0 + h/\sqrt{n}$, for the range of local deviation parameters h specified below.

We assume that the true distribution of covariates X in the training population is known and that the target population shares the same distribution. Specifically, we define F_X as the joint distribution of (a) the truncated normal distribution for the age variable and (b) the binary distribution for the sex variable. For computational purposes, we discretize F_X into 99 bins, each corresponding to a distinct combination of age and sex. Each bin is assigned a probability mass according to F_X , representing the proportion of individuals falling into that bin.

Suppose that the planner has resources to allocate to 75% of the target population. Let

$$W(\theta, \mu) = \int w_R(\theta, x, t, \varepsilon, \lambda) d\mu(x, t).$$

4.1. Finite-sample average optimality. We evaluate performance under a sequence of perturbed DGPs:

$$\theta_{nh} = \theta_0 + h/\sqrt{n}, \quad \text{for } h \in H := \{-2, -1.6, \dots, 2\},$$

where h/\sqrt{n} is added to θ_0 element-wisely. Our goal is to compare the finite-sample average risk

$$\int R(\mu_n^Q, \theta_{nh}) dh = \int \mathbb{E}_{P_{\theta_{nh}}^n} [W_{\mathcal{M}}^*(\theta_{nh}) - W(\theta_{nh}, \mu_n(Z^n))] dh$$

for $Q = P, B$. The simulation proceeds as follows:

- (1) For each h , draw J independent samples of data $\{Z^{n,j}\}_{j=1}^J$ from $P_{\theta_{nh}}^n$ where $Z^{n,j} = \{Z_i^j\}_{i=1}^n$.
- (2) For each j :
 - (a) Obtain the MLE estimates of parameter $\hat{\theta}_{nh}^j$ and Fisher information matrix \hat{I}_{nh}^j .
 - (b) Compute the plug-in rule $\mu_{nh}^{P,j}$ by

$$\mu_{nh}^{P,j} \in \arg \max_{\mu \in \mathcal{M}} W(\hat{\theta}_{nh}^j, \mu).$$

- (c) Draw L samples $\{\theta_\ell\}_{\ell=1}^L$ from $N(\hat{\theta}_{nh}^j, (n\hat{I}_{nh}^j)^{-1})$, which can be interpreted as the quasi-posterior using the quasi-likelihood $N(\hat{\theta}_{nh}^j, (n\hat{I}_{nh}^j)^{-1})$ with the uniform prior (Kim, 2002; Christensen *et al.*, 2026). Then compute the Bayesian rule $\mu_{nh}^{B,j}$ by

$$\mu_{nh}^{B,j} \in \arg \max_{\mu \in \mathcal{M}} \frac{1}{L} \sum_{\ell=1}^L W(\theta_\ell, \mu).$$

- (3) Compute the oracle welfare $W_{\mathcal{M}}^*(\theta_{nh}) = \max_{\mu \in \mathcal{M}} W(\theta_{nh}, \mu)$, and estimate $R(\mu_n^Q(Z^n), \theta_{nh})$ by

$$R(Q, h) := \frac{1}{J} \sum_{j=1}^J [W_{\mathcal{M}}^*(\theta_{nh}) - W(\theta_{nh}, \mu_{nh}^{Q,j})]$$

for $Q = P, B$. Store $R(Q, h)$ for each h .

- (4) Taking the average of $R(Q, h)$ over h gives an estimate of $\int R(\mu_n^Q, \theta_{nh}) dh$ for $Q = P, B$.


We use POT, an open-source Python library developed by Flamary *et al.* (2021), to compute the plug-in rule and the Bayesian rule.

4.2. Results. We study the cases of $n = 200, 500$, $J = 2000$, and $L = 2000$ for both $\lambda = 0$ and $\lambda = 1$. We first report the estimated risks, followed by comparisons of the resulting treatment allocations.

Figure 2 shows the results for $n = 200$. While our theory predicts the plug-in and Bayesian rules are both optimal under smooth welfare ($\lambda = 1$), the simulation shows that the Bayesian rule performs better in small samples. We also observe that the Bayesian rule outperforms the plug-in rule under directionally differentiable welfare ($\lambda = 0$).

Remark 4.1. The finite-sample advantage of the Bayesian rule under smooth welfare ($\lambda = 1$, $n = 200$) is not inconsistent with asymptotic optimality of both rules. Both rules attain the same limit-experiment risk as $n \rightarrow \infty$, but their finite-sample behavior differs because the plug-in rule depends on a point estimate $\hat{\theta}_n$, which can be noisy in small samples. Since the optimal transport plan is sensitive to perturbations of the cost function, small estimation errors in $\hat{\theta}_n$ can lead to large changes in the plug-in allocation. By contrast, the Bayesian rule integrates over the posterior distribution,

FIGURE 2. Comparisons of estimated risks: $n = 200$ (left: smooth welfare, right: directionally differentiable welfare)



effectively averaging over parameter uncertainty and producing a more stable allocation. This robustness advantage diminishes as n increases, as confirmed by the results for $n = 500$ (Figure 3).

Figure 3 shows the results for $n = 500$. The Bayesian rule still performs slightly better when $\lambda = 1$, but the overall risk levels are substantially reduced, and the performance gap between the two rules narrows. This indicates that both rules are approaching optimality as the sample size increases from 200 to 500. When $\lambda = 0$, the Bayesian rule continues to outperform the plug-in rule, which is consistent with our theoretical predictions: under the directionally differentiable welfare the Bayesian rule is optimal, but the plug-in rule may not be. Notably, the Bayesian rule performs particularly well when the values of h are negative. In these cases, the welfare contrasts become smaller, making the assignment problem more challenging. This highlights the robustness of the Bayesian rule to local perturbations that make treatment decisions harder.

To gain further insight into the behavior of the two rules, we visualize the average allocations, $J^{-1} \sum_j \mu_{nh}^{Q,j}$, for $Q = P, E$, under θ_0 , $n = 200$, and $\lambda = 0$. Figure 4 shows that the Bayesian rule deviates from the oracle rule only near the decision boundary, while the plug-in rule exhibits substantial deviations even away from it. This relative stability of the Bayesian rule contributes to a sizable risk reduction. A similar, albeit weaker, pattern is observed for $n = 500$.

Remark 4.2. Under the current simulation setup, when the total amount of resources is sufficiently small, the kink points of the directionally differentiable welfare do not affect the assignment decision, as the available resources are exhausted before the assignment rule encounters the kink points. In such cases, the behavior of the two rules under directionally differentiable welfare resembles their behavior under the smooth welfare. We set the resource level to 75% to allow interaction between the decision boundary under the oracle rule and the kink points.

FIGURE 3. Comparisons of estimated risks: $n = 500$ (left: smooth welfare, right: directionally differentiable welfare)




TABLE 1. Maximum likelihood estimates for math scores

Variables	voucher	age	gender	const	σ^2
Estimates	2.06	-5.5	-0.72	102.77	104.31
Standard error	0.46	0.24	0.44	2.87	5.17

5. EMPIRICAL APPLICATION

We illustrate our methods using data from Angrist *et al.* (2006) continued from Example 2.6.¹² As outcome variables, Angrist *et al.* (2006) use the test scores in language and math. Since the estimation results are similar between these two, we use math scores as the outcome variable for illustration. We focus on the case where the observed test scores are censored at the tenth percentiles of the test score distribution among test-takers (denoted by τ), in line with the original article, to address selection issues. In addition to the test scores, we observe treatment status, as well as age and sex as covariates. The sample includes 3,541 individuals overall, with 1,788 girls and 1,753 boys. Ages range from 10 to 17 with mean 12.7 and standard deviation 1.3. The maximum likelihood estimates are summarized in Table 1.

We then hypothetically treat the marginal distribution of the covariates in the observed sample as that of the target population, and compute both the plug-in and the Bayesian rules as described in the previous section. In this example, the planner's utility function is given by:

$$w_R(\theta, x, t, \varepsilon, \lambda) = \lambda w(\theta, x, t) + (1 - \lambda) \max \{w(\theta, x, t) - \varepsilon, \tau\},$$

where

$$w(\theta, x, t) = (x^\top \beta + \alpha t) + (\tau - x^\top \beta - \alpha t) \Phi \left(\frac{\tau - x^\top \beta - \alpha t}{\sigma} \right) + \sigma \phi \left(\frac{\tau - x^\top \beta - \alpha t}{\sigma} \right).$$

¹²For the replication dataset of the original article, see Angrist *et al.* (2019).

FIGURE 4. Comparison of (average) treatment assignment under directionally differentiable welfare

Note: The upper panels show the oracle rule, the middle show the Bayesian rule, and the lower show the plug-in rule under θ_0 and $n = 200$.

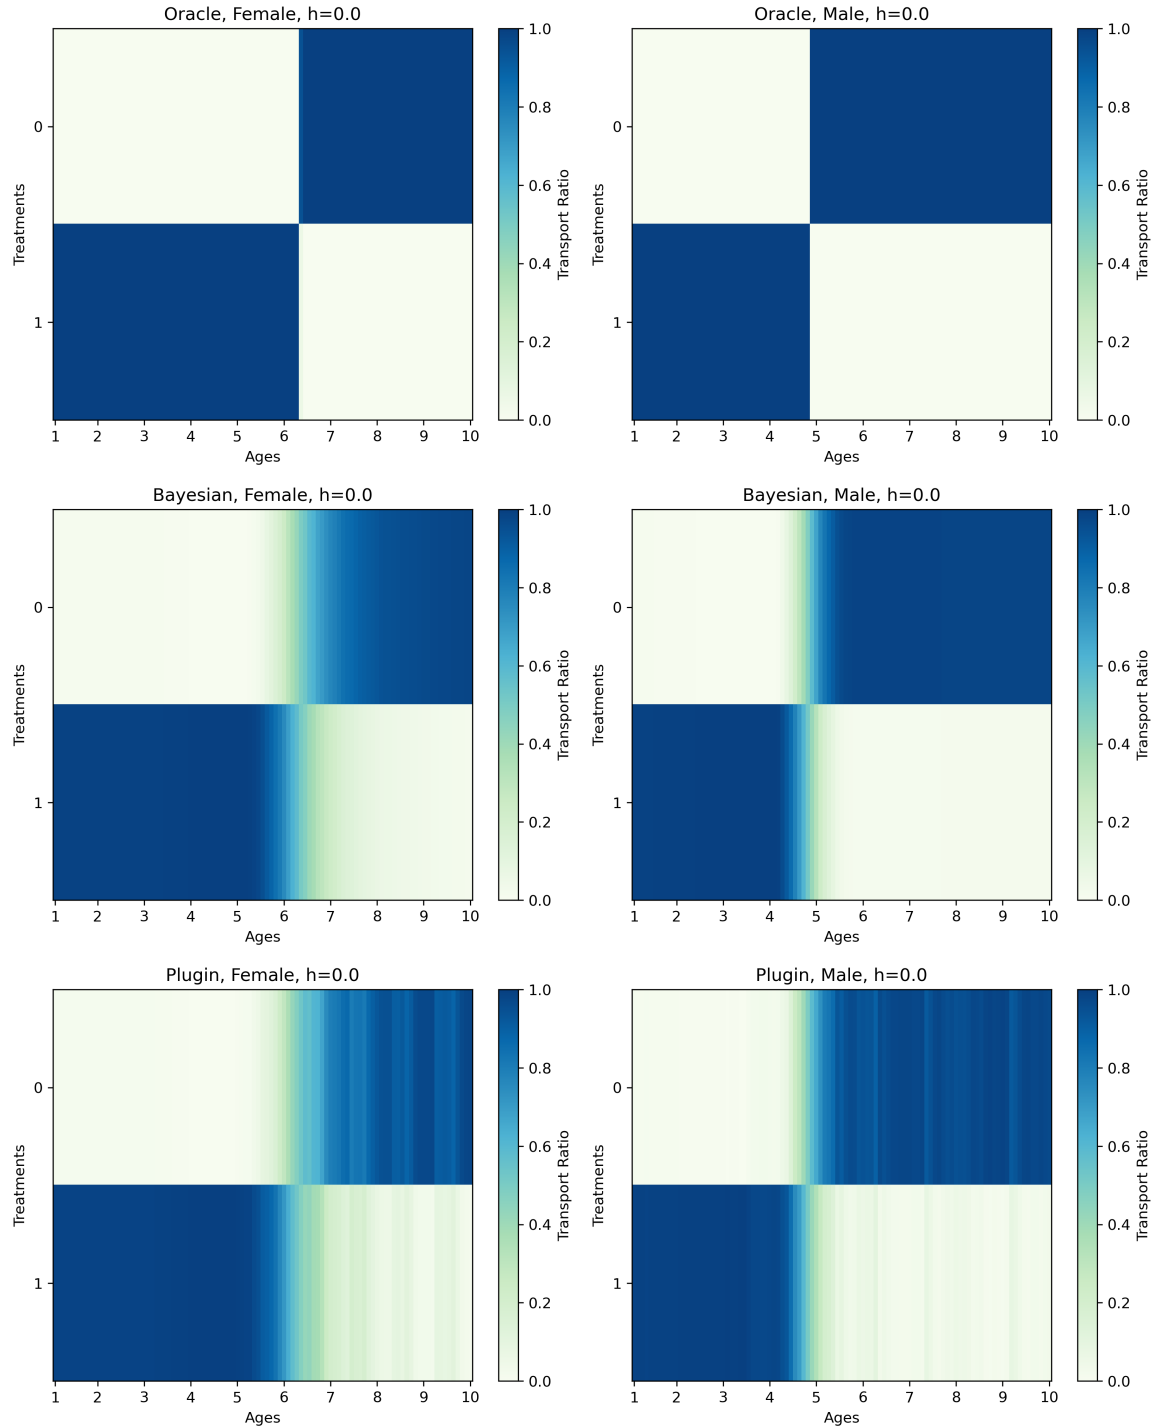
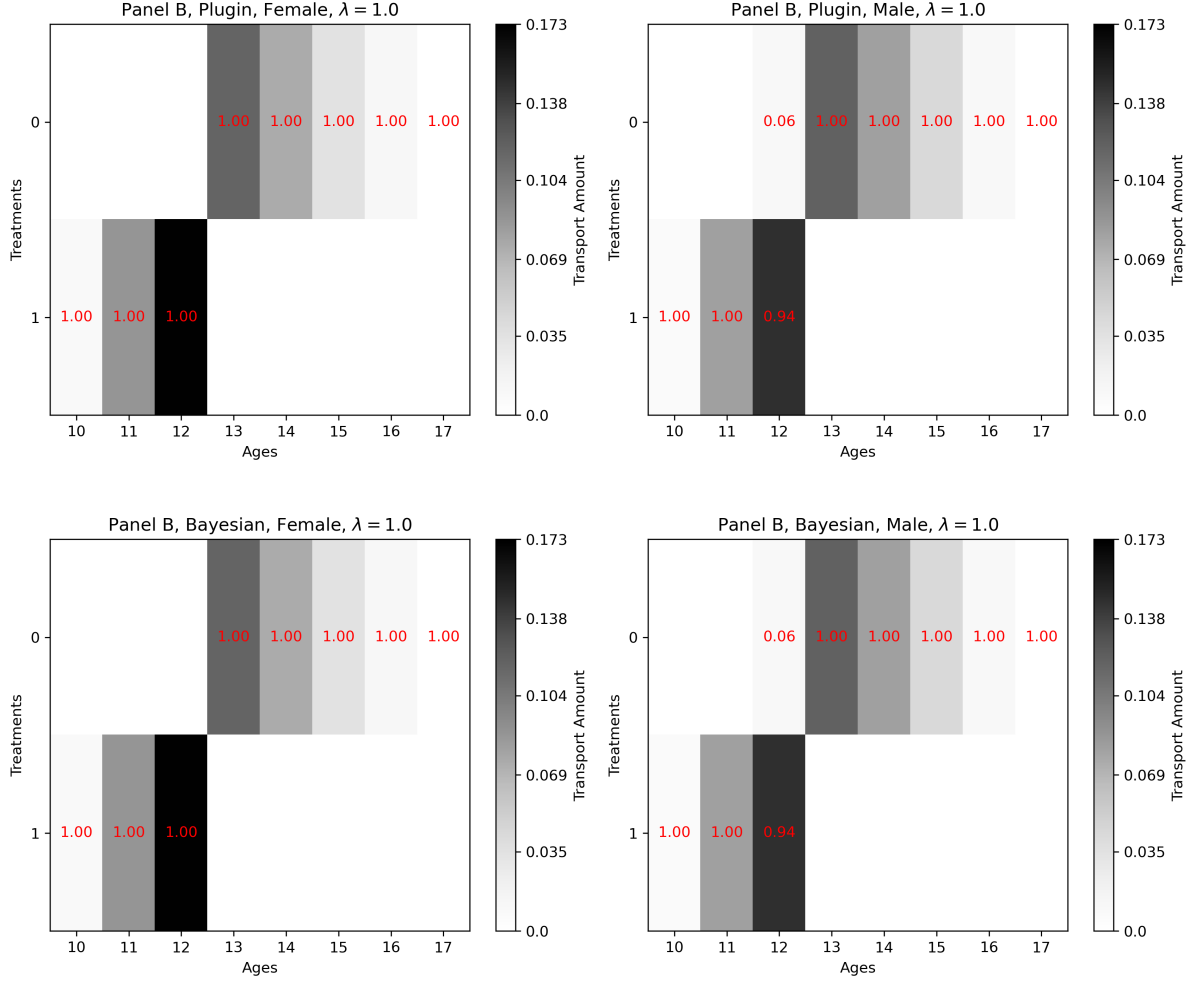


FIGURE 5. Voucher allocations under smooth welfare

Note: The upper panels show the plug-in rule, and the lower panels show the Bayesian rule. The color intensity represents the density of each cell in F_X , with darker shades indicating higher density.

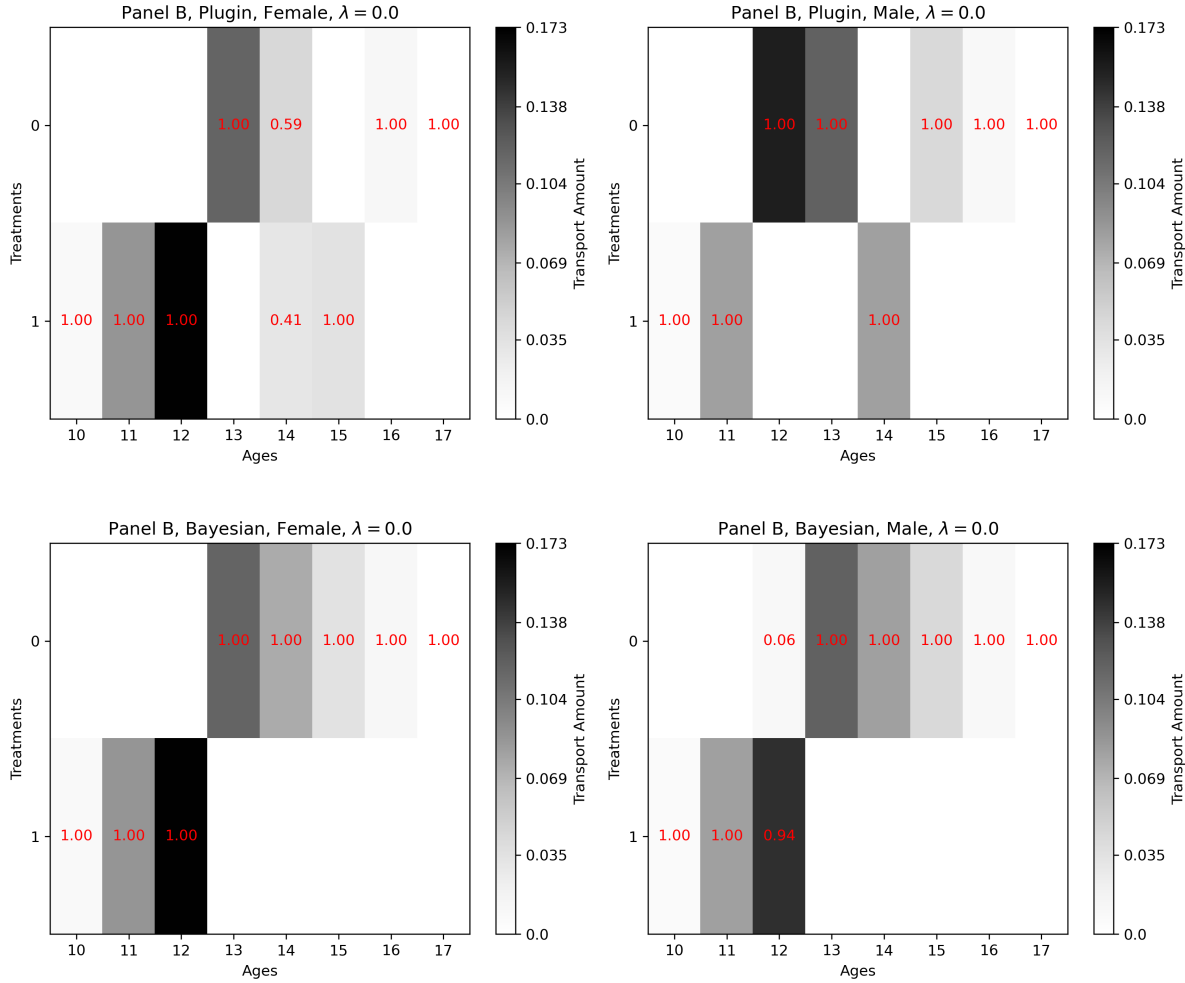


Note that this is slightly different from the utility function in the previous section as the outcome variable is censored at $\tau \neq 0$. In what follows, we focus on $\varepsilon = 3.5$ and $\lambda = 0, 1$. We consider the case where we can assign vouchers for 50% of the target population.

Figure 5 shows the allocations under smooth welfare ($\lambda = 1$). As Table 1 shows, age has a negative effect on outcomes. Accordingly, the plug-in rule allocates vouchers to younger individuals. Since the effect of sex is slightly negative, the plug-in rule prioritizes females over males, resulting in the allocation where vouchers are fully allocated to females aged 10-12, while not fully allocated to males at age 12 as the resource is exhausted due to the capacity constraints. In this setting, the Bayesian rule yields exactly the same allocation, which is natural since both rules are optimal under smooth welfare. This also aligns with the simulation result in the previous section: both rules perform similarly when the sample size is large enough.

FIGURE 6. Voucher allocations under directionally differentiable welfare

Note: The upper panels show the plug-in rule, and the lower panels show the Bayesian rule. The color intensity represents the density of each cell in F_X , with darker shades indicating higher density.



Next, Figure 6 shows the allocations under directionally differentiable welfare ($\lambda = 0$). For the plug-in rule, the value of w_R is censored by τ at age 13 for females and at 12 for males in this setting. As in the previous case, vouchers are fully allocated to females aged 10–12 and males aged 10–11. However, the remaining vouchers are randomly assigned, as the value of w_R is just equal to τ for the rest. For the Bayesian rule, after integrating with respect to the posterior distribution, the value of w_R is censored at τ at age 13 for both females and males. This leads to allocation to males at age 12 until the resource is exhausted, resulting in the same allocation as seen in Figure 5. This illustrates that the plug-in and the Bayesian rules could generate different allocations under $\lambda = 0$.

6. CONCLUSION

We studied the decision-theoretic optimality of treatment assignment rules under capacity constraints on available treatments. Since such constraints complicate the analysis of optimal rules, we transformed the planner’s constrained maximization problem into the unconstrained one using tools from optimal transport theory. This reformulation allows us to search for optimal rules in terms of couplings that automatically satisfy the capacity constraints. We investigated two rules previously studied in the literature—the plug-in rule and the Bayesian rule. Both are optimal in the limit experiment when the planner’s utility function is continuously differentiable; however, the plug-in rule may no longer be optimal when the planner’s utility function is only directionally differentiable. A simulation study supports our theoretical predictions. We demonstrated our methods with a voucher assignment problems for private secondary school attendance using data from Angrist *et al.* (2006).

While we focused on posterior-risk minimization in the Gaussian limit experiment, which is equivalent to the standard average optimality whenever the average risk is well-defined, asymptotic minimax optimality is also a widely used benchmark in local asymptotics frameworks. Kido (2023) provides an asymptotic minimax optimality result when the ATE is partially identified and there are no constraints on available treatments. In that setting, the plug-in rule becomes optimal only when the oracle rule is (Hadamard) differentiable. Given that Hirano and Porter (2009) show the minimax optimality of the plug-in rule under the full differentiability conditions, we conjecture that the modes of differentiability of the planner’s utility function w plays a key role in the minimax optimality of the plug-in rule in our setting.

APPENDIX A. PROOF OF THEOREM 3.4

A.1. Preliminaries. In what follows, the expectation $\mathbb{E}_{P_{\theta_{nh}}^n}$ can be understood as the outer expectation when μ_n is not measurable. Also, P_h denotes the joint law of $\Delta \sim N(h, I_0^{-1})$ and $U \sim \text{Unif}[0, 1]$ in the limit experiment.

Our proof of Theorem 3.4 proceeds as follows. Lemma A.2 shows that the Bayesian rule $\{\mu_n^B\}$ is an element of \mathcal{D} and optimal in the limit experiment. Lemmas A.3–A.7 are used to establish Lemma A.2. Additional auxiliary lemmas are relegated to Appendix B.

The following is a known result and can be found at van der Vaart (1998, Theorem 10.8).

Proposition A.1. *Suppose that model is DQM at θ_0 . Let C_n be the ball of radius M_n for a given, arbitrary sequence $M_n \rightarrow \infty$. Further, suppose $\int \|\theta\|^p d\pi(\theta) < \infty$. Then, for every measurable function f that grows subpolynomially of order p ,*

$$\int f(h) \mathbf{1}_{C_n^c}(h) \pi(\theta_{nh} | Z^n) dh = o_{P_{\theta_0}^n}(1) \quad \text{as } n \rightarrow \infty.$$

A.2. Proof. Let $(\mathbb{D}, \|\cdot\|_{\mathbb{D}})$ be the product metric space induced by (\mathcal{M}, d_W) and $([0, 1], |\cdot|)$. Let $\ell^\infty(\mathbb{D}) := \{f : \mathbb{D} \rightarrow \mathbb{R} : \sup_{(\mu, \varepsilon) \in \mathbb{D}} |f(\mu, \varepsilon)| < \infty\}$. Define

$$\begin{aligned} (\mu, \varepsilon) \mapsto Q_n(\mu, \varepsilon; z) &:= \int \left[\int \sqrt{n} (w(\theta_{nh}, x, t) - w(\theta_0, x, t)) d\mu \right] \pi_n(\theta_{nh} | z) dh - \varepsilon H(\mu), \\ (\mu, \varepsilon) \mapsto Q_\infty(\mu, \varepsilon; \Delta) &:= \int \left[\int \dot{w}_{\theta_0}(x, t; h) d\mu \right] dN(\Delta, I_0^{-1})(h) - \varepsilon H(\mu). \end{aligned}$$

In the same manner as the definition of $\mu_n^B(Z^n)$, we define $\mu_\infty^*(\Delta)$ as the limit of $\mu_{\infty, \varepsilon}^*(\Delta)$ where

$$\mu_{\infty, \varepsilon}^*(\Delta) := \arg \max_{\mu \in A_0} \int \int \dot{w}_{\theta_0}(x, t; s) dN(\Delta, I_0^{-1})(s) d\mu - \varepsilon H(\mu).$$

It is easy to see that $\mu_\infty^*(\Delta) \in \arg \min_{\mu \in A_0} L_\infty(\mu, \Delta)$.

Lemma A.2. *The Bayesian rule $\{\mu_n^B(Z^n)\}$ satisfies (i) $\sqrt{n} P_{\theta_{nh}}^n(\mu_n^B(Z^n) \notin A_0^\delta) \rightarrow 0$ for all $\delta > 0$ and (ii) $\mu_n^B \xrightarrow{h} \mu_\infty^*$ as $n \rightarrow \infty$.*

Proof. (i). We claim that for any $\theta_0 \in \Theta$ and any $\delta > 0$, there are \bar{n} and $\varepsilon'_n(\bar{n})$ (which is at the order of $n^{\alpha+1}$ for some $\alpha \geq 1$) such that for all $n \geq \bar{n}$,

$$P_{\theta_{nh}}^n(\mu_n^B(z) \notin A_0^\delta) \leq P_{\theta_{nh}}^n(\pi_n(N_{1/n}(\theta_0)^c | z) > 2\varepsilon'_n),$$

where $N_\varepsilon(\theta_0) := \{\theta : \|\theta - \theta_0\| < \varepsilon\}$ for $\varepsilon > 0$, and $\pi_n(A | z) := \int_A \pi_n(\theta | z) d\theta$.¹³ Then the conclusion follows since Christensen *et al.* (2026, Lemmas 9–11) imply that $\sqrt{n} P_{\theta_{nh}}^n(\pi_n(N_{1/n}(\theta_0)^c | z) > 2\varepsilon'_n) \rightarrow 0$ as $n \rightarrow \infty$.

It is sufficient to show that for any z ,

$$\pi_n(N_{1/n}(\theta_0)^c | z) \leq 2\varepsilon'_n \implies \mu_n^B(z) \in A_0^\delta. \quad (\text{A.1})$$

¹³This statement is an adaptation of Christensen *et al.* (2026, Lemma 8). Their proof cannot directly apply to our setting because their arguments could fail when the set of actions, \mathcal{M} in our notation, is not finite.

It is trivial if $A_0 = \mathcal{M}$, so suppose $A_0 \subsetneq \mathcal{M}$. By Lemma A.3 below, there exists \bar{n} such that $n \geq \bar{n}$ implies

$$\min_{\mu \in A_0} \int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta|z) d\theta > \sup_{\nu \notin A_0^\delta} \int_{N_{1/n}(\theta_0)} W(\theta, \nu) \pi_n(\theta|z) d\theta$$

Then there exists $\alpha = \alpha(\bar{n}) \geq 1$ such that

$$\min_{\mu \in A_0} \int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta|z) d\theta > \sup_{\nu \notin A_0^\delta} \int_{N_{1/n}(\theta_0)} W(\theta, \nu) \pi_n(\theta|z) d\theta + \frac{1}{\bar{n}^\alpha},$$

which implies that for $n \geq \bar{n}$,

$$\min_{\mu \in A_0} \int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta|z) d\theta > \sup_{\nu \notin A_0^\delta} \int_{N_{1/n}(\theta_0)} W(\theta, \nu) \pi_n(\theta|z) d\theta + \frac{1}{n^\alpha}.$$

Note that we have

$$\begin{aligned} \min_{\mu \in A_0} \left(\int W(\theta, \mu) \pi_n(\theta|z) d\theta \right) &= \min_{\mu \in A_0} \left(\int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta|z) d\theta + \int_{N_{1/n}(\theta_0)^c} W(\theta, \mu) \pi_n(\theta|z) d\theta \right) \\ &\geq \min_{\mu \in A_0} \int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta|z) d\theta + \min_{\mu \in A_0} \int_{N_{1/n}(\theta_0)^c} W(\theta, \mu) \pi_n(\theta|z) d\theta \\ &\geq \min_{\mu \in A_0} \int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta|z) d\theta - \pi_n \left(N_{1/n}(\theta_0)^c | z \right) M, \end{aligned}$$

where $\sup_{\theta, \mu} |W(\theta, \mu)| \leq M < \infty$. Also,

$$\sup_{\nu \notin A_0^\delta} \int W(\theta, \nu) \pi_n(\theta|z) d\theta \leq \sup_{\nu \notin A_0^\delta} \int_{N_{1/n}(\theta_0)} W(\theta, \nu) \pi_n(\theta|z) d\theta + \pi_n \left(N_{1/n}(\theta_0)^c | z \right) M.$$

For the promise for (A.1), we choose $\varepsilon'_n > 0$ that satisfies $(n^{\alpha+1} 2M)^{-1} \leq 2\varepsilon'_n < (n^\alpha 2M)^{-1}$, which leads to

$$\pi_n \left(N_\varepsilon(\theta_0)^c | z \right) < \frac{1}{n^{\alpha+1} 2M}.$$

Then it follows that

$$\begin{aligned} &\min_{\mu \in A_0} \int W(\theta, \mu) \pi_n(\theta|z) d\theta - \sup_{\nu \notin A_0^\delta} \int W(\theta, \nu) \pi_n(\theta|z) d\theta \\ &\geq \min_{\mu \in A_0} \int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta|z) d\theta - \sup_{\nu \notin A_0^\delta} \int_{N_{1/n}(\theta_0)} W(\theta, \nu) \pi_n(\theta|z) d\theta - 2\pi_n \left(N_{1/n}(\theta_0)^c | z \right) M \\ &\geq \frac{1}{n^\alpha} - 2\pi_n \left(N_{1/n}(\theta_0)^c | z \right) M > 0, \end{aligned}$$

which implies

$$\min_{\mu \in A_0} \int W(\theta, \mu) \pi_n(\theta|z) d\theta > \sup_{\nu \notin A_0^\delta} \int W(\theta, \nu) \pi_n(\theta|z) d\theta.$$

Thus we conclude $\mu_n^B(z) \in A_0^\delta$.

(ii). Since $\sqrt{n}P_{\theta_{nh}}^n(\mu_n^B \notin A_0^\delta) \rightarrow 0$ for every $\delta > 0$, the asymptotic distribution of $\mu_n^B(Z^n)$ has support in $\bigcap_{\delta>0} A_0^\delta = A_0$. Hence, for sufficiently large n , μ_n^B equivalently solves

$$\mu_n^B(z) \in \arg \max_{\mu \in A_0} \mathcal{Q}_n(\mu, 0; z).$$

By the definition of μ_n^B ,

$$\mathcal{Q}_n(\mu_n^B(z), 0; z) = \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z).$$

Take any closed subset G of \mathcal{M} . Note that this closedness is in terms of (\mathcal{M}, d_W) . By the Portmanteau lemma, it is sufficient to show that

$$\limsup_{n \rightarrow \infty} P_{\theta_{nh}}^n(\mu_n^B(z) \in G) \leq P_h(\mu_\infty^*(\Delta) \in G)$$

for the conclusion.

By Lemma A.5, for each n ,

$$\{\mu_n^B(z) \in G\} = \left\{ \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon; z) = \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z) \right\}.$$

By Lemmas A.6 and A.7, it follows that $\mathcal{Q}_n \xrightarrow{h} \mathcal{Q}_\infty$ in $\mathcal{F} \subset \ell^\infty(\mathbb{D})$ as $n \rightarrow \infty$, where

$$\mathcal{F} = \left\{ f \in \ell^\infty(\mathbb{D}) : \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} f(\mu, \varepsilon) \text{ exists} \right\}.$$

By Lemma B.3, the operator $f \mapsto \lim_{\varepsilon \downarrow 0} \max_{\mu \in S} f(\mu, \varepsilon)$ is continuous for any closed $S \subset A_0$ at any $f \in \mathcal{F}$. Applying the continuous mapping theorem yields

$$\lim_{\varepsilon \downarrow 0} \max_{\mu \in S} \mathcal{Q}_n(\mu, \varepsilon; z) \xrightarrow{h} \lim_{\varepsilon \downarrow 0} \max_{\mu \in S} \mathcal{Q}_\infty(\mu, \varepsilon; \Delta) \quad \text{as } n \rightarrow \infty.$$

Then

$$\begin{aligned} & \limsup_{n \rightarrow \infty} P_{\theta_{nh}}^n(\mu_n^B(z) \in G) \\ &= \limsup_{n \rightarrow \infty} P_{\theta_{nh}}^n \left(\lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon; z) = \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z) \right) \\ &\leq P_h \left(\lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0 \cap G} \mathcal{Q}_\infty(\mu, \varepsilon; \Delta) = \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_\infty(\mu, \varepsilon; \Delta) \right), \end{aligned}$$

where the inequality follows from the Portmanteau lemma. Since we can derive the equivalence of the events

$$\{\mu_\infty^*(\Delta) \in G\} = \left\{ \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0 \cap G} \mathcal{Q}_\infty(\mu, \varepsilon; \Delta) = \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_\infty(\mu, \varepsilon; \Delta) \right\}$$

in the same manner, applying the Portmanteau lemma again yields $\mu_n^B \xrightarrow{h} \mu_\infty^*$ as $n \rightarrow \infty$. \square

Lemma A.3. *For all $\delta > 0$, there exists $\bar{n} > 0$ such that for all $n \geq \bar{n}$,*

$$\min_{\mu \in A_0} \int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta|z) d\theta > \sup_{\nu \notin A_0^\delta} \int_{N_{1/n}(\theta_0)} W(\theta, \nu) \pi_n(\theta|z) d\theta. \quad (\text{A.2})$$

Proof. Let $g_n(\mu) := \int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta) d\theta$ and $V_{1/n} := \int_{N_{1/n}(\theta_0)} \pi_n(\theta) d\theta$. First, we claim that $V_{1/n}^{-1} g_n(\mu)$ converges to $W(\theta_0, \mu)$ uniformly over \mathcal{M} as $n \rightarrow \infty$; i.e., for all $\eta > 0$, there exists n_η such that for all $\mu \in \mathcal{M}$,

$$n \geq n_\eta \implies \left| V_{1/n}^{-1} g_n(\mu) - W(\theta_0, \mu) \right| < \eta.$$

To show this claim, we argue that (i) for each $\mu \in \mathcal{M}$, $V_{1/n}^{-1} g_n(\mu)$ converges to $W(\theta_0, \mu)$ in pointwise, and (ii) $\{g_n\}_{n \in \mathbb{N}}$ is equicontinuous; i.e., for all $\eta > 0$ and all $\mu \in \mathcal{M}$, there exists $\delta_{(\eta, \mu)} > 0$ such that for all $n \in \mathbb{N}$ and all $\nu \in \mathcal{M}$,

$$d_W(\mu, \nu) < \delta_{(\eta, \mu)} \implies |g_n(\mu) - g_n(\nu)| < \eta.$$

Combining with the compactness of \mathcal{M} , the uniform convergence follows from these two.

To see (i), note that

$$\left| V_{1/n}^{-1} g_n(\mu) - W(\theta_0, \mu) \right| \leq V_{1/n}^{-1} \int_{N_{1/n}(\theta_0)} |W(\theta, \mu) - W(\theta_0, \mu)| \pi_n(\theta|z) d\theta. \quad (\text{A.3})$$

Fix $\eta > 0$. Since the map $\theta \mapsto W(\theta, \mu)$ is continuous at θ_0 , there exists $\delta > 0$ such that

$$\theta \in N_\delta(\theta_0) \implies |W(\theta, \mu) - W(\theta_0, \mu)| < \eta.$$

Then for all n with $n^{-1} < \delta$, RHS of (A.3) is bounded above by η .

To see (ii), fix $\eta > 0$ and $\mu \in \mathcal{M}$. First, note that $g_n(\mu)$ can be written as

$$g_n(\mu) = \mathbb{E}_\mu[\Psi_n(x, t)], \quad \Psi_n(x, t) := \int_{N_{1/n}(\theta_0)} w(\theta, x, t) \pi_n(\theta|z) d\theta,$$

where \mathbb{E}_μ denotes the expectation with respect to the coupling μ . Note that $(x, t) \mapsto w(\theta, x, t)$ is uniformly continuous and bounded since $\mathcal{X} \times \mathcal{T}$ is compact. Then a function $(x, t) \mapsto w_k(\theta, x, t)$ defined by

$$w_k(\theta, x, t) := \inf \{w(\theta, x', t') + kd((x, t), (x', t')) : (x', t') \in \mathcal{X} \times \mathcal{T}\}$$

is k -Lipschitz continuous, and converges uniformly to $w(\theta, x, t)$ from below as $k \rightarrow \infty$ (see e.g., Heinonen (2001, Theorem 6.8)). Lemma A.4 below further extends that the convergence holds uniformly over Θ ; i.e., for all $\eta > 0$, there is a sufficiently large $K = K(\eta)$ such that

$$\sup_{\theta \in \Theta} \max_{(x, t) \in \mathcal{X} \times \mathcal{T}} |w(\theta, x, t) - w_K(\theta, x, t)| < \frac{\eta}{3}.$$

Given this K , define

$$\Psi_n^K(x, t) := \int_{N_{1/n}(\theta_0)} w_K(\theta, x, t) \pi_n(\theta|z) d\theta.$$

Then $\Psi_n^K(x, t)$ is also Lipschitz continuous whose Lipschitz constant is less than or equal to K . Therefore,

$$\begin{aligned} |g_n(\mu) - g_n(\nu)| &\leq \mathbb{E}_\mu \left| \Psi_n(x, t) - \Psi_n^K(x, t) \right| + \left| \mathbb{E}_\mu[\Psi_n^K(x, t)] - \mathbb{E}_\nu[\Psi_n^K(x, t)] \right| + \mathbb{E}_\nu \left| \Psi_n^K(x, t) - \Psi_n(x, t) \right| \\ &< \frac{2}{3} \eta V_{1/n} + K d_W(\mu, \nu) \leq \frac{2}{3} \eta + K d_W(\mu, \nu), \end{aligned}$$

where the second inequality follows from the Kantorovich-Rubinstein duality (Villani, 2009, Theorem 5.10). Thus, we obtain

$$d_W(\mu, \nu) < \frac{\eta}{3K} \implies |g_n(\mu) - g_n(\nu)| < \eta.$$

Therefore $\{g_n\}_{n \in \mathbb{N}}$ is an equicontinuous family.

Finally, we show (A.2). There exists $\eta(\delta) > 0$ such that for all $\mu \in A_0$,

$$W(\theta_0, \mu) > \sup_{\nu \notin A_0^\delta} W(\theta_0, \nu) + \eta(\delta). \quad (\text{A.4})$$

By the uniform convergence shown above, there exists $n_{\eta(\delta)}$ such that for all $\mu \in \mathcal{M}$,

$$n \geq n_{\eta} \implies \left| V_{1/n}^{-1} g_n(\mu) - W(\theta_0, \mu) \right| < \frac{\eta(\delta)}{3}.$$

Then fix $n \geq n_{\eta(\delta)}$, and let $\mu_n \in \arg \min_{\mu \in A_0} g_n(\mu)$. For any $\nu \notin A_0^\delta$, we obtain

$$\begin{aligned} \min_{\mu \in A_0} g_n(\mu) &= g_n(\mu_n) \\ &> V_{1/n} \left[W(\theta_0, \mu_n) - \frac{\eta(\delta)}{3} \right] \\ &> V_{1/n} \left[W(\theta_0, \nu) + \eta(\delta) - \frac{\eta(\delta)}{3} \right] \\ &> g_n(\nu) + V_{1/n} \frac{\eta(\delta)}{3}, \end{aligned}$$

where the second inequality follows from (A.4). Thus we obtain

$$\min_{\mu \in A_0} g_n(\mu) \geq \sup_{\nu \notin A_0^\delta} g_n(\nu) + V_{1/n} \frac{\eta(\delta)}{3}.$$

Hence we conclude $n \geq n_{\eta(\delta)}$ implies $\min_{\mu \in A_0} g_n(\mu) > \sup_{\nu \notin A_0^\delta} g_n(\nu)$. \square

Lemma A.4. *For all $\eta > 0$, there is a sufficiently large $K = K(\eta)$ such that*

$$\sup_{\theta \in \Theta} \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |w(\theta, x, t) - w_K(\theta, x, t)| < \eta.$$

Proof. To simplify the notation, let $\mathcal{Y} = \mathcal{X} \times \mathcal{T}$. Let D be the diameter of \mathcal{Y} . Since $y \mapsto w(\theta, y)$ is continuous uniformly over Θ (Assumption 3.3 (ii)) and \mathcal{Y} is compact (Assumption 3.2), we have

$$\Omega(\delta) := \sup_{\theta \in \Theta} \sup_{d(y, y') < \delta} |w(\theta, y) - w(\theta, y')| \rightarrow 0 \quad \text{as } \delta \downarrow 0.$$

Fix any $\theta \in \Theta$ and $y \in \mathcal{Y}$. Then for all $y' \in \mathcal{Y}$,

$$w(\theta, y') \geq w(\theta, y) - |w(\theta, y) - w(\theta, y')| \geq w(\theta, y) - \Omega(d(y, y')),$$

where the last inequality follows from the definition of Ω . Adding $kd(y, y')$ to both sides and taking the infimum with respect to y' yields

$$w_k(\theta, y) \geq w(\theta, y) + \inf_{r \in [0, D]} \{kr - \Omega(r)\},$$

which implies

$$w(\theta, y) - w_k(\theta, y) \leq \sup_{r \in [0, D]} \phi_k(r), \quad \phi_k(r) := \Omega(r) - kr.$$

By the definition of Ω , there exists $\delta > 0$ such that $\Omega(\delta) < \eta$. Note that Ω is non-decreasing function. Then if $0 \leq r < \delta$, we have $\phi_k(r) \leq \Omega(\delta) < \eta$. If $r \geq \delta$, we have $\phi_k(r) \leq \Omega(D) - k\delta$. Thus $\sup_{r \in [0, D]} \phi_k(r) \leq \{\Omega(\delta), \Omega(D) - k\delta\}$. Hence, for sufficiently large K , it follows $\sup_{r \in [0, D]} \phi_K(r) < \eta$. This implies

$$\sup_{\theta \in \Theta} \max_{y \in \mathcal{Y}} \{w(\theta, y) - w_K(\theta, y)\} \leq \sup_{r \in [0, D]} \phi_K(r) < \eta.$$

Note that it always holds $w_k(\theta, y) \leq w(\theta, y)$ for each k . Thus we conclude the proof. \square

Lemma A.5. *For each n and each closed $G \subset \mathcal{M}$, the following equivalence of the events holds:*

$$\left\{ \mu_n^B(z) \in G \right\} = \left\{ \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon; z) = \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z) \right\}.$$

Proof. (C). Take any $z \in \mathcal{Z}^n$ such that $\left\{ \mu_n^B(z) \in G \right\}$. Note that

$$\lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon; z) \leq \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z)$$

is clear. Suppose by way of contradiction that

$$\lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon; z) < \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z).$$

Then, there is small enough $\varepsilon_1 > 0$ such that

$$\max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon_1; z) < \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon_1; z).$$

We can take $\eta > 0$ such that

$$\max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon_1; z) + 2\eta < \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon_1; z)$$

Since $\max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z) \rightarrow \max_{\mu \in A_0} \mathcal{Q}_n(\mu, 0; z)$ as $\varepsilon \downarrow 0$ from Lemmas F.5 and F.6, there exists small enough $\varepsilon_2 > 0$ such that

$$\left| \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon_2; z) - \max_{\mu \in A_0} \mathcal{Q}_n(\mu, 0; z) \right| < \eta.$$

Also, because $\varepsilon H(\mu_n^B) \rightarrow 0$ as $\varepsilon \downarrow 0$, there exists small enough $\varepsilon_3 > 0$ such that $\varepsilon_3 H(\mu_n^B) < \eta$. Let $\varepsilon = \min \{\varepsilon_1, \varepsilon_2, \varepsilon_3\}$. Then,

$$\begin{aligned} \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z) &< \max_{\mu \in A_0} \mathcal{Q}_n(\mu, 0; z) + \eta \\ &= \int \int \sqrt{n} (w(\theta_{nh}, x, t) - w(\theta_0, x, t)) d\mu_n^B(z) d\pi(\theta_{nh}|z) + \eta \\ &\leq \int \int \sqrt{n} (w(\theta_{nh}, x, t) - w(\theta_0, x, t)) d\mu_n^B(z) d\pi(\theta_{nh}|z) - \varepsilon H(\mu_n^B) + \varepsilon H(\mu_n^B) + \eta \\ &\leq \max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon; z) + \varepsilon H(\mu_n^B) + \eta \\ &< \max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon; z) + 2\eta < \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z), \end{aligned}$$

which is a contradiction. Hence it holds $\lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon; z) = \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z)$.

(\supset). For the other direction, take any $z \in \mathcal{Z}^n$ with

$$\lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon; z) = \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon; z).$$

Let $\varepsilon_k \downarrow 0$ as $k \rightarrow \infty$. Recall

$$\mu_{n, \varepsilon_k}^B(z) = \arg \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon_k; z).$$

By Proposition 3.3, $\mu_{n, \varepsilon_k}^B(z)$ converges to $\mu_n^B(z)$ weakly as $k \rightarrow \infty$. Moreover, by Villani (2009, Theorem 6.9), we know that weak convergence in \mathcal{M} is equivalent to convergence in (\mathcal{M}, d_W) . Hence, $\mu_{n, \varepsilon_k}^B(z) \rightarrow \mu_n^B(z)$ in the Wasserstein distance d_W .

First, we argue that for each $K \in \mathbb{N}$ there is $k \geq K$ such that $\mu_{n, \varepsilon_k}^B(z) \in G$. By way of contradiction, assume that there is K such that for any $k \geq K$, $\mu_{n, \varepsilon_k}^B(z) \notin G$. Now, $\mu_{n, \varepsilon_k}^B(z) \notin G$ implies that

$$\max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon_k; z) < \mathcal{Q}_n(\mu_{n, \varepsilon_k}^B, \varepsilon_k; z).$$

Take any $\eta > 0$ such that

$$\max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon_k; z) + \eta < \mathcal{Q}_n(\mu_{n, \varepsilon_k}^B, \varepsilon_k; z).$$

Because

$$\lim_{k \rightarrow \infty} \left(\max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon_k; z) - \max_{\mu \in (A_0 \cap G)} \mathcal{Q}_n(\mu, \varepsilon_k; z) \right) = 0,$$

there exists K' such that if $k \geq K'$ then

$$\left| \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon_k; z) - \max_{\mu \in (A_0 \cap G)} \mathcal{Q}_n(\mu, \varepsilon_k; z) \right| = \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon_k; z) - \max_{\mu \in (A_0 \cap G)} \mathcal{Q}_n(\mu, \varepsilon_k; z) < \eta/2.$$

Therefore, for sufficiently large k , we have

$$\max_{\mu \in A_0 \cap G} \mathcal{Q}_n(\mu, \varepsilon_k; z) + \eta < \mathcal{Q}_n(\mu_{n, \varepsilon_k}^B, \varepsilon_k; z) \leq \max_{\mu \in A_0} \mathcal{Q}_n(\mu, \varepsilon_k; z) < \max_{\mu \in (A_0 \cap G)} \mathcal{Q}_n(\mu, \varepsilon_k; z) - \eta/2,$$

which leads to a contradiction. Therefore, for each K there is $k \geq K$ such that $\mu_{n, \varepsilon_k}^B(z) \in G$.

Now, create such a subsequence $\{\mu_{n,\varepsilon_{k_\ell}}^B(z)\}_{\ell \in \mathbb{N}}$ with $\mu_{n,\varepsilon_{k_\ell}}^B(z) \in G$ for each ℓ . Note that any subsequence of convergent sequence in arbitrary metric space converges to the same limit as the original sequence. Therefore, $\mu_{n,\varepsilon_{k_\ell}}^B(z) \rightarrow \mu_n^B(z)$ in d_W as $\ell \rightarrow \infty$. Since G is closed, we conclude that $\mu_n^B(z) \in G$. \square

Let

$$Q_n(\mu, \varepsilon; z) = \int \left(\int \dot{w}_{\theta_0}(x, t; h) d\mu \right) \pi_n(\theta_{nh}|z) dh - \varepsilon H(\mu),$$

$$\tilde{Q}_n(\mu, \varepsilon; z) = \int \left(\int \dot{w}_{\theta_0}(x, t; h) d\mu \right) dN(\Delta_n(z), I_0^{-1})(h) - \varepsilon H(\mu),$$

where $\Delta_n(z) = \frac{1}{\sqrt{n}} \sum_{i=1}^n I_0^{-1} s(z_i) \in \mathbb{R}^k$ with the score function s at θ_0 and $\Delta_n \xrightarrow{0} G \sim N(0, I_0^{-1})$. Define $\mathbb{D} \equiv \mathcal{M} \times [0, 1]$. Let $\{Q_n(\mu, \varepsilon) : (\mu, \varepsilon) \in \mathbb{D}\}$, $\{\tilde{Q}_n(\mu, \varepsilon) : (\mu, \varepsilon) \in \mathbb{D}\}$, $\{Q_\infty(\mu, \varepsilon) : (\mu, \varepsilon) \in \mathbb{D}\}$ be stochastic processes. Assume that they yield maps $Q_n : \mathcal{Z}^n \rightarrow \ell^\infty(\mathbb{D})$, $\tilde{Q}_n : \mathcal{Z}^n \rightarrow \ell^\infty(\mathbb{D})$, and $Q_\infty : \mathbb{R}^k \rightarrow \ell^\infty(\mathbb{D})$. We can do this since the sample paths are continuous by Lemma B.2.

Let

$$\mathcal{F} = \left\{ f \in \ell^\infty(\mathbb{D}) : \lim_{\varepsilon \downarrow 0} \max_{\mu \in A_0} f(\mu, \varepsilon) \text{ exists} \right\},$$

where the sup-norm is equipped to \mathcal{F} . Note that $Q_n(z), \tilde{Q}_n(z), Q_\infty(\Delta) \in \mathcal{F}$ for any z and Δ .

Lemma A.6. $Q_n \xrightarrow{h} Q_\infty$ in \mathcal{F} as $n \rightarrow \infty$.

Proof. Note that $Q_n = Q_n + o_{P_{\theta_{nh}}}^n(1)$ as a process in \mathcal{F} as $n \rightarrow \infty$ by Lemma A.7. Hence, it is sufficient to show that $Q_n \xrightarrow{h} Q_\infty$ in \mathcal{F} as $n \rightarrow \infty$. Let $C_M \subset \mathbb{R}^k$ be the closed ball of radius M around 0. Define stochastic processes by

$$Q_{n,M}(\mu, \varepsilon; z) = \int_{C_M} \left(\int \dot{w}_{\theta_0}(x, t; h) d\mu(x, t) \right) \pi_n(\theta_{nh}|z) dh - \varepsilon H(\mu),$$

$$\tilde{Q}_{n,M}(\mu, \varepsilon; z) = \int_{C_M} \left(\int \dot{w}_{\theta_0}(x, t; h) d\mu(x, t) \right) dN(\Delta_n(z), I_0^{-1})(h) - \varepsilon H(\mu),$$

$$Q_{\infty,M}(\mu, \varepsilon; \Delta) = \int_{C_M} \left(\int \dot{w}_{\theta_0}(x, t; h) d\mu(x, t) \right) dN(\Delta, I_0^{-1})(h) - \varepsilon H(\mu).$$

First, we will show that $Q_{n,M} - \tilde{Q}_{n,M} \xrightarrow{h} 0$ in \mathcal{F} as $n \rightarrow \infty$ for any fixed M . Note that

$$\begin{aligned} & \left\| Q_{n,M}(\cdot; Z^n) - \tilde{Q}_{n,M}(\cdot; Z^n) \right\|_{\mathcal{F}} \\ & \leq \max_{h \in C_M} \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\dot{w}_{\theta_0}(x, t; h)| \cdot \left\| \pi_n(\theta_{nh}|Z^n) - N(\Delta_n(Z^n), I_0^{-1}) \right\|_{\text{TV}} \end{aligned}$$

where $\|\cdot\|_{\text{TV}}$ is the total variation norm. Because C_M is bounded and $\dot{w}_{\theta_0}(x, t; h)$ is bounded for any h , we know that $\max_{h \in C_M} \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\dot{w}_{\theta_0}(x, t; h)| < \infty$. From the Bernstein–von Mises theorem,

$$\left\| \pi_n(\theta_{nh}|Z^n) - N(\Delta_n(Z^n), I_0^{-1}) \right\|_{\text{TV}} \xrightarrow{0} 0 \quad \text{as } n \rightarrow \infty.$$

From Le Cam's first lemma, it also converges to 0 along $P_{\theta_{nh}}^n$. Therefore, $Q_{n,M} - \tilde{Q}_{n,M} \xrightarrow{h} 0$ in \mathcal{F} as $n \rightarrow \infty$.

Next, we argue that $\tilde{Q}_{n,M} \xrightarrow{h} Q_M$ in \mathcal{F} as $n \rightarrow \infty$ for any fixed M . Define $\phi : \mathbb{R}^k \rightarrow \mathcal{F}$ by

$$\phi(\delta)(\mu, \varepsilon) = \int_{C_M} \left(\int \dot{w}_{\theta_0}(x, t; h) d\mu(x, t) \right) dN(\delta, I_0^{-1})(h) - \varepsilon H(\mu)$$

Since ϕ is continuous, and $\Delta_n \xrightarrow{h} \Delta \sim N(h, I_0^{-1})$ as $n \rightarrow \infty$ by Le Cam's third lemma, the continuous mapping theorem implies $\phi(\Delta_n) \xrightarrow{h} \phi(\Delta)$ as $n \rightarrow \infty$. Thus $\tilde{Q}_{n,M} \xrightarrow{h} Q_M$ in \mathcal{F} as $n \rightarrow \infty$.

Combining the above two findings, we obtain $Q_{n,M} \xrightarrow{h} Q_{\infty,M}$ in \mathcal{F} as $n \rightarrow \infty$ from the Slutsky theorem. We also have that $Q_{\infty,M} - Q_{\infty} = o_{P_h^\Delta}(1)$ as $M \rightarrow \infty$ where P_h^Δ is the (marginal) law of $\Delta \sim N(h, I_0^{-1})$. Thus, there exists a sequence $M_n \rightarrow \infty$ such that $Q_{n,M_n} \xrightarrow{h} Q_{\infty}$ in \mathcal{F} as $n \rightarrow \infty$.

Finally, it remains to show that $Q_n - Q_{n,M_n} = o_{P_{\theta_{nh}}^n}(1)$ in \mathcal{F} as $n \rightarrow \infty$, which leads to the conclusion, $Q_n \xrightarrow{h} Q_{\infty}$ in \mathcal{F} as $n \rightarrow \infty$. By Assumption 3.4 (iii),

$$\left| \int_{\mathbb{R}^k \setminus C_{M_n}} \left(\int \dot{w}_{\theta_0}(x, t; h) d\mu(x, t) \right) \pi_n(\theta_{nh}|z) dh \right| \leq \int_{\mathbb{R}^k \setminus C_{M_n}} K(h) \pi_n(\theta_{nh}|z) dh.$$

Then applying Proposition A.1 yields that RHS is $o_{P_{\theta_0}^n}(1)$. Thus it is $o_{P_{\theta_{nh}}^n}(1)$ as well. Hence it follows that $Q_n(\mu, \varepsilon; Z^n) - Q_{n,M_n}(\mu, \varepsilon; Z^n) \xrightarrow{h} 0$ for any $(\mu, \varepsilon) \in \mathbb{D}$. Then the continuity of sample path implies $Q_n - Q_{n,M_n} = o_{P_{\theta_{nh}}^n}(1)$ in \mathcal{F} as desired. \square

Lemma A.7. $Q_n = Q_{\infty} + o_{P_{\theta_{nh}}^n}(1)$ in \mathcal{F} as $n \rightarrow \infty$.

Proof. Notice that for every $z \in \mathcal{Z}^n$, $\sup_{(\mu, \varepsilon) \in \mathbb{D}} |Q_n(\mu, \varepsilon; z) - Q_{\infty}(\mu, \varepsilon; z)|$ is bounded above by

$$\int \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\tilde{w}_{n,0}(x, t; h)| \pi_n(\theta_{nh}|z) dh,$$

where $\tilde{w}_{n,0}(x, t; h) := \sqrt{n}(w(\theta_{nh}, x, t) - w(\theta_0, x, t)) - \dot{w}_{\theta_0}(x, t; h)$. Let C_{M_n} be a closed ball of radius M_n around 0, where M_n is the divergent sequence specified in Lemma B.4. Then the previous display is further bounded by

$$\pi_n(\theta_{nh}|z)(C_{M_n}) \times \max_{h \in C_{M_n}} \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\tilde{w}_{n,0}(x, t; h)| + \int \mathbf{1}_{\mathbb{R}^k \setminus C_{M_n}}(h) \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\tilde{w}_{n,0}(x, t; h)| \pi_n(\theta_{nh}|z) dh \quad (\text{A.5})$$

The first term of (A.5) converges to zero as $n \rightarrow \infty$ by Lemma B.4. For the second term of (A.5), note that from Assumption 3.4 (i),

$$\max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\sqrt{n}(w(\theta_{nh}, x, t) - w(\theta_0, x, t))| \leq \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\dot{w}_{\theta_0}(x, t; h)| + o(1),$$

From Assumption 3.4 (iii), $\max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\dot{w}_{\theta_0}(x, t; h)|$ is bounded by $K(h)$ that grows at subpolynomially of order p . This implies that $\max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |\tilde{w}_{n,0}(x, t; h)|$ is also dominated by a function that grows subpolynomially of order p for sufficiently large n . Then applying Proposition A.1 yields the conclusion. \square

APPENDIX B. AUXILIARY LEMMAS FOR THEOREM 3.4

The following extends the result on Hadamard directional differentiability given by Römisch (2004, Proposition 1). Compared to his setting, the objective map $\int w(\theta, x, t) d\mu$ need not to be linear in θ . By leveraging the uniform continuity from Assumptions 3.3 and 3.4, we obtain the similar form of the directional derivative as his result.

Lemma B.1. *For a closed set $S \subset \mathcal{M}$, define the map $W_S^* : \Theta \rightarrow \mathbb{R}$ by*

$$W_S^*(\theta) = \max_{\mu \in S} \int w(\theta, x, t) d\mu.$$

Then W_S^ is Hadamard directionally differentiable with derivative*

$$\dot{W}_{S,0}^*[h] \equiv \lim_{\varepsilon \downarrow 0} \sup_{\mu \in S^\varepsilon(\theta_0)} \int \dot{w}_{\theta_0}(x, t; h) d\mu$$

where

$$S^\varepsilon(\theta) \equiv \left\{ \mu \in S : \int w(\theta, x, t) d\mu + \varepsilon \geq \max_{\mu \in S} \int w(\theta, x, t) d\mu \right\} \neq \emptyset$$

for $\varepsilon > 0$ and $\theta \in \Theta$. Moreover, if $S \subset A_0$ then $\dot{W}_{S,0}^*[h] = \max_{\mu \in S} \int \dot{w}_{\theta_0}(x, t; h) d\mu$.

Proof. The second statement follows from the first statement and the fact that $S^\varepsilon(\theta_0) = S$ for any $\varepsilon > 0$. Hereafter, we focus on the first statement.

Fix any closed $S \subset \mathcal{M}$ and any $\theta_0 \in \Theta$. Then, $\arg \max_{\mu \in S} W(\theta_0, \mu) \neq \emptyset$ because $w(\theta_0, \cdot)$ is bounded continuous on $\mathcal{X} \times \mathcal{T}$ from Assumption 3.3 (Villani, 2009, Theorem 4.1). Let $\mu(\theta_0) \in \arg \max_{\mu \in S} W(\theta_0, \mu)$. Since $\arg \max_{\mu \in S} W(\theta_0, \mu) \subset S^\varepsilon(\theta_0)$ for any $\varepsilon > 0$, we can guarantee $S^\varepsilon(\theta_0) \neq \emptyset$. Also, because $w(\theta_0, \cdot, \cdot)$ is bounded, $\max_{\mu \in S} \int w(\theta_0, x, t) d\mu < \infty$.

Let $r_n \downarrow 0$, and $h_n \rightarrow h$. Define

$$\sigma_n = \frac{1}{r_n} (W_S^*(\theta_0 + r_n h_n) - W_S^*(\theta_0))$$

and we want to show that $\sigma_n \rightarrow \dot{W}_{S,0}^*[h]$. But, it is enough to show that for any subsequence of $\{\sigma_n\}$, there exists a further subsequence that converges to $\dot{W}_{S,0}^*[h]$. Take any subsequence and denote it by $\{\sigma_n\}$ for simplicity.

Define the maps $T_n : S \rightarrow \mathbb{R}$ and $T : S \rightarrow \mathbb{R}$ by

$$\begin{aligned} T_n(\mu) &= \int \frac{1}{r_n} (w(\theta_0 + r_n h_n, x, t) - w(\theta_0, x, t)) d\mu, \\ T(\mu) &= \int \dot{w}_{\theta_0}(x, t; h) d\mu. \end{aligned}$$

First, to see that $T_n \rightarrow T$ uniformly, take any $\mu \in S$. Note that

$$|T_n(\mu) - T(\mu)| \leq \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} \left| \frac{1}{r_n} (w(\theta_0 + r_n h_n, x, t) - w(\theta_0, x, t)) - \dot{w}_{\theta_0}(x, t; h) \right|.$$

We can make RHS arbitrary small as $n \rightarrow \infty$ without depending on μ from Assumption 3.4 (i). Hence, we conclude that $T_n \rightarrow T$ uniformly.

Thus, for each n there exists K_1 such that for each $k \geq K_1$

$$\left| \int \frac{1}{r_{n_k}} (w(\theta_0 + r_{n_k} h_{n_k}, x, t) - w(\theta_0, x, t)) d\mu - \int \dot{w}_{\theta_0}(x, t; h) d\mu \right| < r_n^2 \quad \forall \mu \in S. \quad (\text{B.1})$$

Moreover, from Assumption 3.3 (i), there exists K_2 such that for each $k \geq K_2$

$$\left| \int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\mu - \int w(\theta_0, x, t) d\mu \right| < r_n^2/2 \quad \forall \mu \in S. \quad (\text{B.2})$$

Let $k \geq \max\{K_1, K_2\}$, and construct a further subsequence $\{\sigma_{n_k}\}$.

For each k , take any $\mu_{n_k} \in S^{r_n^2 r_{n_k}}(\theta_0)$. Then, from $\mu_{n_k} \in S$ and the definition of $S^{r_n^2 r_{n_k}}(\theta_0)$,

$$\begin{aligned} \frac{1}{r_{n_k}} (W_S^*(\theta_0 + r_{n_k} h_{n_k}) - W_S^*(\theta_0)) &= \frac{1}{r_{n_k}} \left(\max_{\mu \in S} \int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\mu - \max_{\mu \in S} \int w(\theta_0, x, t) d\mu \right) \\ &\geq \frac{1}{r_{n_k}} \left(\int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\mu_{n_k} - \max_{\mu \in S} \int w(\theta_0, x, t) d\mu \right) \\ &\geq \frac{1}{r_{n_k}} \left(\int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\mu_{n_k} - \int w(\theta_0, x, t) d\mu_{n_k} \right) - r_n^2. \end{aligned}$$

From (B.1), we have

$$\sigma_{n_k} > \int \dot{w}_0(x, t; h) d\mu_{n_k} - 2r_n^2.$$

Therefore,

$$\sigma_{n_k} \geq \sup_{\mu \in S^{r_n^2 r_{n_k}}(\theta_0)} \int \dot{w}_0(x, t; h) d\mu - 2r_n^2.$$

Since $r_n^2 \downarrow 0$, it leads to

$$\liminf_{k \rightarrow \infty} \sigma_{n_k} \geq \lim_{\varepsilon \downarrow 0} \sup_{\mu \in S^\varepsilon(\theta_0)} \int \dot{w}_0(x, t; h) d\mu.$$

Also, take any $\mu'_{n_k} \in S^{r_n^2 r_{n_k}}(\theta_0 + r_{n_k} h_{n_k})$. Then,

$$\begin{aligned} \frac{1}{r_{n_k}} (W_S^*(\theta_0 + r_{n_k} h_{n_k}) - W_S^*(\theta_0)) &= \frac{1}{r_{n_k}} \left(\max_{\mu \in S} \int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\mu - \max_{\mu \in S} \int w(\theta_0, x, t) d\mu \right) \\ &\leq \frac{1}{r_{n_k}} \left(\int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\mu'_{n_k} - \int w(\theta_0, x, t) d\mu'_{n_k} \right) + r_n^2. \end{aligned}$$

From (B.1),

$$\sigma_{n_k} < \int \dot{w}_{\theta_0}(x, t; h) d\mu'_{n_k} + 2r_n^2.$$

If we would have $S^{r_n^2 r_{n_k}}(\theta_0 + r_{n_k} h_{n_k}) \subset S^{r_n^2 r_{n_k} + r_n^2}(\theta_0)$, then we obtain

$$\sigma_{n_k} \leq \sup \left\{ \int \dot{w}_{\theta_0}(x, t; h) d\mu : \mu \in S^{r_n^2 r_{n_k} + r_n^2}(\theta_0) \right\} + 2r_n^2,$$

which leads to

$$\limsup_{k \rightarrow \infty} \sigma_{n_k} \leq \lim_{\varepsilon \downarrow 0} \sup_{\mu \in S^\varepsilon(\theta_0)} \int \dot{w}_{\theta_0}(x, t; h) d\mu,$$

thus $\sigma_{n_k} \rightarrow \dot{W}_{S,0}^*[h]$. Hence, it suffices to show $S^{r_n^2 r_{n_k}}(\theta_0 + r_{n_k} h_{n_k}) \subset S^{r_n^2 r_{n_k} + r_n^2}(\theta_0)$ for the conclusion. Take any $\nu \in S^{r_n^2 r_{n_k}}(\theta_0 + r_{n_k} h_{n_k})$, then

$$\int w(\theta_0 + r_{n_k} h_{n_k}) d\nu + r_n^2 r_{n_k} \geq \max_{\mu \in S} \int w(\theta_0 + r_{n_k} h_{n_k}) d\mu. \quad (\text{B.3})$$

Let $\mu(\theta_0) \in \arg \max_{\mu \in S} \int w(\theta_0, x, t) d\mu$, then

$$\begin{aligned} & \int w(\theta_0, x, t) d\nu + r_n^2 r_{n_k} + r_n^2 \\ \geq & \int w(\theta_0, x, t) d\nu + r_n^2 + \max_{\mu \in S} \int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\mu - \int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\nu \quad \because (\text{B.3}) \\ \geq & \int w(\theta_0, x, t) d\nu + r_n^2 + \int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\mu(\theta_0) - \int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\nu \\ = & \int w(\theta_0, x, t) d\nu + r_n^2 \\ & + \int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\mu(\theta_0) - \int w(\theta_0, x, t) d\mu(\theta_0) + \int w(\theta_0, x, t) d\mu(\theta_0) - \int w(\theta_0 + r_{n_k} h_{n_k}, x, t) d\nu \\ = & \int w(\theta_0, x, t) d\mu(\theta_0) + r_n^2 \\ & + \int (w(\theta_0, x, t) - w(\theta_0 + r_{n_k} h_{n_k}, x, t)) d\nu + \int (w(\theta_0 + r_{n_k} h_{n_k}, x, t) - w(\theta_0, x, t)) d\mu(\theta_0) \\ \geq & \int w(\theta_0, x, t) d\mu(\theta_0) + r_n^2 - r_n^2/2 - r_n^2/2 \quad \because (\text{B.2}) \\ = & \int w(\theta_0, x, t) d\mu(\theta_0). \end{aligned}$$

Thus, $\nu \in S^{r_n^2 r_{n_k} + r_n^2}(\theta_0)$. □

Lemma B.2. *The sample paths of Q_n , \tilde{Q}_n , and Q_∞ are continuous and bounded in \mathcal{M} .*

Proof. First, we will show that the sample path of

$$Q_n(\mu, \varepsilon; z) = \int \left(\int \dot{w}_{\theta_0}(x, t; h) d\mu \right) \pi_n(\theta_{nh}|z) dh - \varepsilon H(\mu)$$

is continuous. Note that, then, it is bounded because \mathcal{M} is compact. Fix any z , and take any $\{(\mu_k, \varepsilon_k)\}_{k=1}^\infty \subset \mathbb{D}$ that converges to (μ, ε) . Since \mathbb{D} is a metric space, overall convergence implies elementwise convergence. Thus, $\mu_k \rightarrow \mu$ in the Wasserstein distance and $\varepsilon_k \rightarrow \varepsilon$. We are done if

$$\left| \int \left\{ \int \dot{w}_{\theta_0}(x, t; h) d\mu_k - \int \dot{w}_{\theta_0}(x, t; h) d\mu \right\} \pi_n(\theta_{nh}|z) dh - \varepsilon_k H(\mu_k) + \varepsilon H(\mu) \right| \rightarrow 0,$$

as $k \rightarrow \infty$.

By the triangle inequality, LHS is bounded above by

$$\left| \int \left\{ \int \dot{w}_{\theta_0}(x, t; h) d\mu_k - \int \dot{w}_{\theta_0}(x, t; h) d\mu \right\} \pi_n(\theta_{nh}|z) dh \right| + |\varepsilon_k H(\mu_k) - \varepsilon H(\mu)|.$$

The second term converges to zero since H is continuous. For the first term, since $\mu_k \rightarrow \mu$ in the Wasserstein distance implies $\mu_k \rightsquigarrow \mu$, we have

$$\begin{aligned} \int \left(\int \dot{w}_{\theta_0}(x, t; h) d\mu_k \right) \pi_n(\theta_{nh}|z) dh &= \int \left(\int \dot{w}_{\theta_0}(x, t; h) \pi_n(\theta_{nh}|z) dh \right) d\mu_k \\ &\rightarrow \int \left(\int \dot{w}_{\theta_0}(x, t; h) \pi_n(\theta_{nh}|z) dh \right) d\mu, \end{aligned}$$

where the convergence holds if the map $(x, t) \mapsto \int \dot{w}_{\theta_0}(x, t; h) \pi_n(\theta_{nh}|z) dh$ is continuous and bounded. To see the continuity, let $\{(x_k, t_k)\}$ be such that $(x_k, t_k) \rightarrow (x, t)$. Then

$$\begin{aligned} &\left| \int \dot{w}_{\theta_0}(x, t; h) \pi_n(\theta_{nh}|z) dh - \int \dot{w}_{\theta_0}(x_k, t_k; h) \pi_n(\theta_{nh}|z) dh \right| \\ &\leq \int |\dot{w}_{\theta_0}(x, t; h) - \dot{w}_{\theta_0}(x_k, t_k; h)| \pi_n(\theta_{nh}|z) dh \rightarrow 0 \quad \text{as } k \rightarrow \infty, \end{aligned}$$

where the convergence follows from the continuity of $\dot{w}_{\theta_0}(x, t; h)$ in (x, t) and the dominated convergence theorem. The boundedness of $(x, t) \mapsto \int \dot{w}_{\theta_0}(x, t; h) \pi_n(\theta_{nh}|z) dh$ follows from the continuity of $\dot{w}_{\theta_0}(x, t; h)$ in (x, t) (Assumption 3.4 (ii)) and the compactness of $\mathcal{X} \times \mathcal{T}$ (Assumption 3.2).

Similar arguments can be applied to \tilde{Q}_n and Q_∞ . \square

Lemma B.3. *For any $S \subset A_0$. the operator $M : \mathcal{F} \rightarrow \mathbb{R}$ where $M(f) := \lim_{\varepsilon \downarrow 0} \max_{\mu \in S} f(\mu, \varepsilon)$ is continuous.*

Proof. Let $f_k \rightarrow f$ in \mathcal{F} as $k \rightarrow \infty$. Note that $M(f)$ and $M(f_k)$ exist for each k by the definition of \mathcal{F} . Then

$$\begin{aligned} |M(f) - M(f_k)| &= \lim_{\varepsilon \downarrow 0} \left| \max_{\mu \in \mathcal{M}} f(\mu, \varepsilon) - \max_{\mu \in \mathcal{M}} f_k(\mu, \varepsilon) \right| \\ &\leq \lim_{\varepsilon \downarrow 0} \max_{\mu \in \mathcal{M}} |f(\mu, \varepsilon) - f_k(\mu, \varepsilon)| \\ &\leq \max_{(\mu, \varepsilon) \in \mathbb{D}} |f(\mu, \varepsilon) - f_k(\mu, \varepsilon)| \rightarrow 0, \end{aligned}$$

as $k \rightarrow \infty$, where the equality follows because $|\cdot|$ is continuous, and the convergence follows because $f_k \rightarrow f$ in \mathcal{F} as $k \rightarrow \infty$. \square

The next result is an adaptation of Christensen *et al.* (2026, Lemma 3), where we need a modification to allow max-operator within the expression.

Lemma B.4. *There is a sequence $\{M_n\}$ such that $M_n \uparrow \infty$, $M_n/\sqrt{n} \rightarrow 0$, and*

$$\sup_{\|h\| \leq 2M_n} \max_{(x, t) \in \mathcal{X} \times \mathcal{T}} |\sqrt{n} [w(\theta_{nh}, x, t) - w(\theta_0, x, t)] - \dot{w}_{\theta_0}(x, t; h)| \rightarrow 0.$$

Proof. From Shapiro (1990, Lemmas 3.3 and 3.4) and Assumption 3.4 (i), we know that for any compact $S \subset \mathbb{R}^k$ and $\varepsilon > 0$, there is N such that $\sup_{h \in S} g_n(h) < \varepsilon$ for any $n \geq N$ where $g_n(h) = \max_{(x, t) \in \mathcal{X} \times \mathcal{T}} |\sqrt{n} [w(\theta_{nh}, x, t) - w(\theta_0, x, t)] - \dot{w}_{\theta_0}(x, t; h)|$. Define

$$\psi_n = \sup_{\|h\| \leq 2 \log(1+n)} g_n(h).$$

We are done if $\psi_n \rightarrow 0$ because $\log(1+n) \uparrow \infty$ and $n^{-1/2} \log(n+1) \rightarrow 0$. To show $\psi_n \rightarrow 0$, it is enough to show that for any subsequence ψ_n (abusing notation) there is a further subsequence converging to 0. First, consider $\sup_{\|h\| \leq 2 \log(1+1)} g_n(h)$. We know that there is $N(1)$ such that $\sup_{\|h\| \leq 2 \log(1+1)} g_n(h) < 1/\log(1+1)$ for any $n \geq N(1)$. Second, for $\sup_{\|h\| \leq 2 \log(1+2)} g_n(h)$, there is $N(2)$ such that $\sup_{\|h\| \leq 2 \log(1+2)} g_n(h) < 1/\log(1+2)$ for any $n \geq N(2)$. Proceed with $N(1) < N(2) < N(3) < \dots$ WLOG. Then, the map $N : \mathbb{N} \rightarrow \mathbb{N}$ satisfies $i < j \implies N(i) < N(j)$. Hence, $\psi_{N(n)}$ is a subsequence of ψ_n . To show $\psi_{N(n)} = \sup_{\|h\| \leq 2 \log(1+n)} g_{N(n)}(h) \rightarrow 0$, take any $\varepsilon > 0$. Since $1/\log(1+n) \downarrow 0$, there is \bar{n} such that $1/\log(1+n) < \varepsilon$ for any $n \geq \bar{n}$. Therefore, $\psi_{N(n)} < 1/\log(1+n) < \varepsilon$ for any $n \geq \bar{n}$. \square

APPENDIX C. PROOF OF THEOREM 3.10

Let $\{\mu_n\} \in \mathcal{D}$ be any sequence of decision rules that will be matched by μ_∞ in the limit experiment. We show the followings: (i)

$$\int \liminf_{n \rightarrow \infty} \sqrt{n} R(\mu_n, \theta_{nh}) dh \geq \int_0^1 \int L_\infty(\mu_\infty(\Delta, u), \Delta) d\Delta du,$$

(ii)

$$\int \limsup_{n \rightarrow \infty} \sqrt{n} R(\mu_n, \theta_{nh}) dh \leq \int_0^1 \int L_\infty(\mu_\infty(\Delta, u), \Delta) d\Delta du,$$

and (iii) $\{\mu_n\}$ is average optimal if the matched rule μ_∞ of μ_n satisfies $\mu_\infty(\Delta) \in \arg \min_{\mu \in A_0} L_\infty(\mu, \Delta)$ for every realization of $\Delta \sim N(h, I_0^{-1})$.

(i). For any h ,

$$\begin{aligned} & \mathbb{E}_{P_{\theta_{nh}}^n} \sqrt{n} [W_{\mathcal{M}}^*(\theta_{nh}) - W(\theta_{nh}, \mu_n(Z^n))] \\ &= \int \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_{nh}, x, t) d\mu - \max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu \right) dP_{\theta_{nh}}^n(z) \\ & \quad + \int \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu - \int w(\theta_0, x, t) d\mu_n(z) \right) dP_{\theta_{nh}}^n(z) \\ & \quad - \int \sqrt{n} \left(\int w(\theta_{nh}, x, t) d\mu_n(z) - \int w(\theta_0, x, t) d\mu_n(z) \right) dP_{\theta_{nh}}^n(z). \end{aligned} \tag{C.1}$$

For the first term of the RHS of (C.1),

$$\begin{aligned} & \int \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_{nh}, x, t) d\mu - \max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu \right) dP_{\theta_{nh}}^n(z) \\ &= \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_{nh}, x, t) d\mu - \max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu \right) \rightarrow \dot{W}_{\mathcal{M}, 0}^*[h], \end{aligned}$$

where the convergence follows from Lemma B.1. For the second term of the RHS of (C.1), it is clear that

$$\int \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu - \int w(\theta_0, x, t) d\mu_n(z) \right) dP_{\theta_{nh}}^n(z) \geq 0.$$

For the third term of the RHS of (C.1),

$$\begin{aligned} & \int \sqrt{n} \left(\int w(\theta_{nh}, x, t) d\mu_n(z) - \int w(\theta_0, x, t) d\mu_n(z) \right) dP_{\theta_{nh}}^n(z) \\ &= \int \int \dot{w}_{\theta_0}(x, t; h) d\mu_n(z) dP_{\theta_{nh}}^n(z) \\ & \quad + \int \int \{ \sqrt{n} [w(\theta_{nh}, x, t) - w(\theta_0, x, t)] - \dot{w}_{\theta_0}(x, t; h) \} d\mu_n(z) dP_{\theta_{nh}}^n(z) \end{aligned}$$

where $\dot{w}_{\theta_0}(x, t; h)$ is the directional derivative of $w(\theta, x, t)$ at θ_0 . By Assumption 3.4 (i),

$$\begin{aligned} & \left| \int \int \{ \sqrt{n} [w(\theta_{nh}, x, t) - w(\theta_0, x, t)] - \dot{w}_{\theta_0}(x, t; h) \} d\mu_n(z) dP_{\theta_{nh}}^n(z) \right| \\ & \leq \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} | \sqrt{n} [w(\theta_{nh}, x, t) - w(\theta_0, x, t)] - \dot{w}_{\theta_0}(x, t; h) | \rightarrow 0. \end{aligned}$$

Therefore,

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \mathbb{E}_{P_{\theta_{nh}}^n} \sqrt{n} [W_{\mathcal{M}}^*(\theta_{nh}) - W(\theta_{nh}, \mu_n(Z^n))] \\ & \geq \dot{W}_{\mathcal{M},0}^*[h] - \limsup_{n \rightarrow \infty} \int \int \dot{w}_{\theta_0}(x, t; h) d\mu_n(z) dP_{\theta_{nh}}^n(z) \\ & \equiv \dot{W}_{\mathcal{M},0}^*[h] - \limsup_{n \rightarrow \infty} \int \varphi(\mu_n(z)) dP_{\theta_{nh}}^n(z) \\ & \geq \dot{W}_{\mathcal{M},0}^*[h] - \mathbb{E}_{(\Delta, U) \sim P_h} \varphi(\mu_\infty(\Delta, U)) \end{aligned}$$

from the portmanteau theorem in metric spaces since the map $\varphi : \mu \mapsto \int \dot{w}_{\theta_0}(x, t; h) d\mu$ is bounded continuous on \mathcal{M} . Thus, it follows that

$$\begin{aligned} & \int \liminf_{n \rightarrow \infty} \mathbb{E}_{P_{\theta_{nh}}^n} \sqrt{n} [W_{\mathcal{M}}^*(\theta_{nh}) - W(\theta_{nh}, \mu_n(Z^n))] dh \\ & \geq \int \mathbb{E}_{P_h} [\dot{W}_{\mathcal{M},0}^*[h] - \varphi(\mu_\infty)] dh \\ & = \int \int_0^1 \int_{\mathbb{R}^k} \left[\dot{W}_{\mathcal{M},0}^*[h] - \int \dot{w}_{\theta_0}(x, t; s) d\mu_\infty(\Delta, u) \right] dN(h, I_0^{-1})(\Delta) du dh \\ & = \int_0^1 \int_{\mathbb{R}^k} \int \left[\dot{W}_{\mathcal{M},0}^*[s] - \int \dot{w}_{\theta_0}(x, t; s) d\mu_\infty(\Delta, u) \right] dN(\Delta, I_0^{-1})(s) d\Delta du, \end{aligned}$$

where the last equality follows by Tonelli's theorem since the integrand is nonnegative. By the definition of L_∞ , the last display is equal to

$$\int_0^1 \int L_\infty(\mu_\infty(\Delta, u), \Delta) d\Delta du.$$

It should be noted that $L_\infty(\mu_\infty(\Delta, u), \Delta)$ can depend on u only through μ_∞ .

(ii) The argument can be carried out analogously to (i). However, for the second term of the RHS of (C.1),

$$\begin{aligned}
& \int \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu - \int w(\theta_0, x, t) d\mu_n(z) \right) dP_{\theta_{nh}}^n(z) \\
&= \int_{\{\mu_n \in A_0\}} \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu - \int w(\theta_0, x, t) d\mu_n(z) \right) dP_{\theta_{nh}}^n(z) \\
&\quad + \int_{\{\mu_n \notin A_0\}} \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu - \int w(\theta_0, x, t) d\mu_n(z) \right) dP_{\theta_{nh}}^n(z) \\
&= \int_{\{\mu_n \notin A_0\}} \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu - \int w(\theta_0, x, t) d\mu_n(z) \right) dP_{\theta_{nh}}^n(z),
\end{aligned}$$

where the second equality follows from $\max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu = \max_{\mu \in A_0} \int w(\theta_0, x, t) d\mu$. Further,

$$\begin{aligned}
& \int_{\{\mu_n \notin A_0\}} \sqrt{n} \left(\max_{\mu \in \mathcal{M}} \int w(\theta_0, x, t) d\mu - \int w(\theta_0, x, t) d\mu_n(z) \right) dP_{\theta_{nh}}^n(z) \\
&\leq \sqrt{n} P_{\theta_{nh}}^n(\mu_n \notin A_0) \times 2 \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |w(\theta_0, x, t)| \rightarrow 0,
\end{aligned}$$

as $n \rightarrow \infty$ from the definition of $\{\mu_n\} \in \mathcal{D}$. Thus, applying the portmanteau theorem yields

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{P_{\theta_{nh}}^n} \sqrt{n} [W_{\mathcal{M}}^*(\theta_{nh}) - W(\theta_{nh}, \mu_n(Z^n))] \leq \dot{W}_{\mathcal{M},0}^*[h] - \mathbb{E}_{P_h} \left[\int \dot{w}_{\theta_0}(x, t; h) d\mu_{\infty} \right]. \quad (\text{C.2})$$

Then we can conclude similarly as in (i).

(iii). Fix any sequence of rules $\{\mu'_n\} \in \mathcal{D}$, and let μ'_{∞} be the matched rule of μ'_n . Combining (i) and (ii) yields

$$\begin{aligned}
\int \liminf_{n \rightarrow \infty} \sqrt{n} R(\mu_n, \theta_{nh}) dh &\leq \int \limsup_{n \rightarrow \infty} \sqrt{n} R(\mu_n, \theta_{nh}) dh \\
&\leq \int_0^1 \int L_{\infty}(\mu_{\infty}(\Delta, u), \Delta) d\Delta du \\
&\leq \int_0^1 \int L_{\infty}(\mu'_{\infty}(\Delta, u), \Delta) d\Delta du \\
&\leq \int \liminf_{n \rightarrow \infty} \sqrt{n} R(\mu'_n, \theta_{nh}) dh,
\end{aligned}$$

which completes the proof.

APPENDIX D. PROOF OF COROLLARY 3.12

The proof of Theorem 3.10 in Appendix C shows that average optimality follows once we establish three statements: (i) $\int \liminf_{n \rightarrow \infty} \sqrt{n} R(\mu_n^B, \theta_{nh}) dh \geq \int_0^1 \int L_{\infty}(\mu_{\infty}^*(\Delta, u), \Delta) d\Delta du$ (lower bound on risk), (ii) $\int \limsup_{n \rightarrow \infty} \sqrt{n} R(\mu_n^B, \theta_{nh}) dh \leq \int_0^1 \int L_{\infty}(\mu_{\infty}^*(\Delta, u), \Delta) d\Delta du$ (upper bound on risk), and (iii) μ_{∞}^* minimizes L_{∞} over A_0 . Statement (i) holds for any $\{\mu_n\} \in \mathcal{D}$ without the additional condition $\sqrt{n} P_{\theta_{nh}}^n(\mu_n \notin A_0) \rightarrow 0$. Statement (iii) follows from Theorem 3.4. The additional condition is used only in the proof of statement (ii), where it is needed to show that

$$\mathbb{E}_{P_{\theta_{nh}}^n} \left[\sqrt{n} \left(W_{\mathcal{M}}^*(\theta_0) - W(\theta_0, \mu_n^B(Z^n)) \right) \right] \rightarrow 0. \quad (\text{D.1})$$

We show that (D.1) holds under Assumption 3.6 and the gap condition $\eta(\delta) \geq c\delta^\gamma$ with $\gamma < 2\alpha$, without imposing the additional condition.

Choose $\delta_n = n^{-\beta}$ with $\beta \in (1/2, \alpha/\gamma)$. Such β exists since $\gamma < 2\alpha$ implies $\alpha/\gamma > 1/2$.

Step 1: $\mu_n^B \in A_0^{\delta_n}$ with high probability. By Assumption 3.6 (ii), for any $\theta \in N_{1/n}(\theta_0)$ and any $\mu \in \mathcal{M}$,

$$|W(\theta, \mu) - W(\theta_0, \mu)| \leq \max_{(x,t) \in \mathcal{X} \times \mathcal{T}} |w(\theta, x, t) - w(\theta_0, x, t)| \leq C \|\theta - \theta_0\|^\alpha \leq Cn^{-\alpha}.$$

It follows that the uniform convergence in the proof of Lemma A.3 has the rate

$$\sup_{\mu \in \mathcal{M}} \left| V_{1/n}^{-1} g_n(\mu) - W(\theta_0, \mu) \right| \leq Cn^{-\alpha},$$

where $g_n(\mu) := \int_{N_{1/n}(\theta_0)} W(\theta, \mu) \pi_n(\theta|z) d\theta$ and $V_{1/n} := \int_{N_{1/n}(\theta_0)} \pi_n(\theta|z) d\theta$. We now apply the argument of Lemma A.3 with $\delta = \delta_n$. Since $\eta(\delta_n) \geq c n^{-\beta\gamma}$ and $\beta\gamma < \alpha$, for sufficiently large n we have $\eta(\delta_n)/3 > Cn^{-\alpha}$, and thus

$$\min_{\mu \in A_0} g_n(\mu) > \sup_{\nu \notin A_0^{\delta_n}} g_n(\nu) + \frac{\eta(\delta_n)}{3} V_{1/n}.$$

The same decomposition as in the proof of Lemma A.2 (i), combined with the posterior concentration from Christensen *et al.* (2026, Lemmas 9–11), then gives

$$\sqrt{n} P_{\theta_{nh}}^n (\mu_n^B \notin A_0^{\delta_n}) \rightarrow 0. \quad (\text{D.2})$$

Step 2: Bounding the welfare gap. By Assumption 3.6 (i) and the Kantorovich–Rubinstein duality (Villani, 2009, Theorem 5.10), $|W(\theta, \mu) - W(\theta, \nu)| \leq L d_W(\mu, \nu)$ for all $\theta \in \Theta$ and $\mu, \nu \in \mathcal{M}$. On the event $\{\mu_n^B(z) \in A_0^{\delta_n}\}$, there exists $\mu_0 \in A_0$ with $d_W(\mu_n^B(z), \mu_0) \leq \delta_n$ (by the definition of $A_0^{\delta_n}$ and the compactness of A_0). Since $\mu_0 \in A_0$, we have $W(\theta_0, \mu_0) = W_{\mathcal{M}}^*(\theta_0)$, and thus

$$W_{\mathcal{M}}^*(\theta_0) - W(\theta_0, \mu_n^B(z)) = W(\theta_0, \mu_0) - W(\theta_0, \mu_n^B(z)) \leq L d_W(\mu_0, \mu_n^B(z)) \leq L \delta_n.$$

On the complementary event $\{\mu_n^B(z) \notin A_0^{\delta_n}\}$, we use the crude bound $W_{\mathcal{M}}^*(\theta_0) - W(\theta_0, \mu_n^B(z)) \leq 2M$, where $M := \sup_{\theta, \mu} |W(\theta, \mu)| < \infty$. Combining the two cases:

$$\begin{aligned} \mathbb{E}_{P_{\theta_{nh}}^n} \left[\sqrt{n} \left(W_{\mathcal{M}}^*(\theta_0) - W(\theta_0, \mu_n^B(z)) \right) \right] &\leq L \sqrt{n} \delta_n + 2M \sqrt{n} P_{\theta_{nh}}^n (\mu_n^B \notin A_0^{\delta_n}) \\ &= Ln^{1/2-\beta} + o(1) \rightarrow 0, \end{aligned}$$

where the second term is $o(1)$ by (D.2) and the convergence $Ln^{1/2-\beta} \rightarrow 0$ follows from $\beta > 1/2$.

Step 3: Conclusion. Since (D.1) holds, the proof of statement (ii) in Appendix C goes through for the Bayesian rule $\{\mu_n^B\}$. Combined with statements (i) and (iii), we conclude that the Bayesian rule is average optimal whenever the average optimality is well-defined.

APPENDIX E. PROOF OF PROPOSITION 3.2

By Villani (2009, Corollary 6.11), $\mu \mapsto d_W(\mu, \nu)$ is continuous on \mathcal{M} .¹⁴ Convexity of $\mu \mapsto d_W(\mu, \nu)$ is shown below. Because the map $\mu \mapsto d_W(\mu, \nu)$ is convex and nonnegative and the map $\mathbb{R}_+ \ni x \mapsto x^2$ is increasing and strictly convex, the composite map $\mu \mapsto (d_W(\mu, \nu))^2$ is strictly convex and nonnegative. Also, it is bounded by the compactness of \mathcal{M} and the continuity.

To see the convexity of $\mu \mapsto d_W(\mu, \nu)$, fix any $\mu_1, \mu_2 \in \mathcal{M}$ and $\alpha \in (0, 1)$. To simplify the notation, let $\mathcal{Y} = \mathcal{X} \times \mathcal{T}$. Let

$$\begin{aligned}\gamma_1^* &\in \arg \inf_{\gamma \in \Gamma(\mu_1, \nu)} \int d(y, y') \gamma(dy, dy'), \\ \gamma_2^* &\in \arg \inf_{\gamma \in \Gamma(\mu_2, \nu)} \int d(y, y') \gamma(dy, dy').\end{aligned}$$

First, we show that $\alpha\gamma_1^* + (1 - \alpha)\gamma_2^* \in \Gamma(\alpha\mu_1 + (1 - \alpha)\mu_2, \nu)$. Take any $A, B \in \mathcal{Y}$. Then,

$$(\alpha\gamma_1^* + (1 - \alpha)\gamma_2^*)(\mathcal{Y} \times B) = \alpha\gamma_1^*(\mathcal{Y} \times B) + (1 - \alpha)\gamma_2^*(\mathcal{Y} \times B) = \nu(B).$$

Also,

$$\begin{aligned} &(\alpha\gamma_1^* + (1 - \alpha)\gamma_2^*)(A \times \mathcal{Y}) \\ &= \alpha\gamma_1^*(A \times \mathcal{Y}) + (1 - \alpha)\gamma_2^*(A \times \mathcal{Y}) = \alpha\mu_1(A) + (1 - \alpha)\mu_2(A) = (\alpha\mu_1 + (1 - \alpha)\mu_2)(A).\end{aligned}$$

Hence,

$$\begin{aligned} r(\alpha\mu_1 + (1 - \alpha)\mu_2) &\leq \int d(y, y') d(\alpha\gamma_1^* + (1 - \alpha)\gamma_2^*) \\ &= \alpha \int d(y, y') d\gamma_1^* + (1 - \alpha) \int d(y, y') d\gamma_2^* = \alpha r(\mu_1) + (1 - \alpha)r(\mu_2).\end{aligned}$$

APPENDIX F. PROOF OF PROPOSITION 3.3

We provide a proof under a general framework using continuous and bounded cost function $c : \mathcal{X} \times \mathcal{T} \rightarrow \mathbb{R}$. Define

$$\begin{aligned}\mathcal{C}_\varepsilon &:= \inf_{\mu \in \mathcal{M}} \int cd\mu + \varepsilon H(\mu). \quad (\varepsilon\text{EOT}) \\ \mathcal{C}_0 &:= \inf_{\mu \in \mathcal{M}} \int cd\mu. \quad (\text{OT}).\end{aligned}$$

Let $\mathcal{M}_{opt} = \arg \min_{\mu \in \mathcal{M}} \int cd\mu$. It should be noted that (\mathcal{M}, d_W) is a metric space which is convex and compact.

Our proof of Proposition 3.3 proceeds as follows. Lemma F.4 gives the conclusion under the assumption $\lim_{\varepsilon \rightarrow 0} \mathcal{C}_\varepsilon = \mathcal{C}_0$. Lemmas F.1–F.3 are needed to prove Lemma F.4. Finally, Lemmas F.5 and F.6 show $\lim_{\varepsilon \rightarrow 0} \mathcal{C}_\varepsilon = \mathcal{C}_0$.

¹⁴More explicitly, if μ_k converges to μ weakly in \mathcal{M} , then $d_W(\mu_k, \nu) \rightarrow d_W(\mu, \nu)$.

Lemma F.1. *Let $\mu_n \in \mathcal{M}$. Suppose that $\lim_n H(\mu_n) =: a \in \mathbb{R}$ exists and that*

$$\limsup_{m,n \rightarrow \infty} H(\mu_{m,n}) \geq a$$

for $\mu_{m,n} := (\mu_m + \mu_n)/2$. Then $\{\mu_n\}$ converges weakly.

Proof. Let $D < \infty$ be the diameter of $\mathcal{X} \times \mathcal{T}$. By Villani (2009, Theorem 6.15), $d_W(\mu, \nu)$ is bounded by $D \|\mu - \nu\|_{\text{TV}}$. By following the same arguments of Nutz (2022, Lemma 1.9), we obtain

$$\lim_{m,n \rightarrow \infty} \|\mu_m - \mu_n\|_{\text{TV}} = 0.$$

Thus it follows that $\lim_{m,n \rightarrow \infty} d_W(\mu_m, \mu_n) = 0$. \square

The next result is an adaptation of Nutz (2022, Theorem 1.10). We use weak convergence as the mode of convergence, whereas the original proof uses convergence in total variation.

Lemma F.2. *Let $\mathcal{Q} \subset \mathcal{M}$ be a convex and closed subset. There exists a unique $\mu_* \in \mathcal{Q}$ such that*

$$H(\mu_*) = \inf_{\mu \in \mathcal{Q}} H(\mu) \in [0, \infty).$$

Proof. Let $\mu_n \in \mathcal{Q}$ be such that $H(\mu_n) \rightarrow \inf_{\mu' \in \mathcal{Q}} H(\mu')$. By convexity of \mathcal{Q} , we have $\mu_{m,n} := (\mu_m + \mu_n)/2 \in \mathcal{Q}$ and hence $H(\mu_{m,n}) \geq \inf_{\mu \in \mathcal{Q}} H(\mu)$ for all m, n . Lemma F.1 shows that $\{\mu_n\}$ converges weakly to some μ_* . By the continuity of $\mu \mapsto H(\mu)$, μ_* is a minimizer of $\inf_{\mu' \in \mathcal{Q}} H(\mu')$. Uniqueness follows from the strict convexity of H . \square

The next result is an adaptation of Nutz (2022, Proposition 1.17).

Lemma F.3. *Consider a decreasing sequence of sets $\mathcal{Q}_n \subset \mathcal{M}$ that are convex and closed, and let $\mathcal{Q} := \bigcap_n \mathcal{Q}_n$. Let $\mu_n = \arg \min_{\mu \in \mathcal{Q}_n} H(\mu)$ be the minimizer of \mathcal{Q}_n . Then*

$$\mu_n \rightarrow \mu_* \text{ weakly, and } H(\mu_n) \rightarrow H(\mu_*),$$

where $\mu_* = \arg \min_{\mu' \in \mathcal{Q}} H(\mu')$.

Proof. Note that the inclusion $\mathcal{Q}_n \supset \mathcal{Q}_{n+1} \supset \mathcal{Q}$ implies that $H(\mu_n)$ is increasing and $H(\mu_n) \leq \inf_{\mu' \in \mathcal{Q}} H(\mu')$. Since any increasing and bounded-above sequence is convergent, we have $\lim H(\mu_n) \leq \inf_{\mu' \in \mathcal{Q}} H(\mu') < \infty$. For $m \geq n$, we have $\mu_{m,n} := (\mu_m + \mu_n)/2 \in \mathcal{Q}_n$ by convexity. Then $H(\mu_{m,n}) \geq H(\mu_n)$. Thus $\limsup_{m,n \rightarrow \infty} H(\mu_{m,n}) \geq \lim H(\mu_n)$. Since $\lim H(\mu_n) < \infty$, Lemma F.1 implies that μ_n converges weakly to some limit μ . By the continuity of H on \mathcal{M} ,

$$H(\mu) = \lim_n H(\mu_n) \leq \inf_{\mu' \in \mathcal{Q}} H(\mu').$$

Thus we obtain $\mu \in \arg \min_{\mu' \in \mathcal{Q}} H(\mu')$. By the uniqueness of the minimizer shown in Lemma F.2, we have $\mu = \mu_*$; i.e., μ_n converges weakly to μ_* . \square

The next result is an adaptation of Nutz (2022, Theorem 5.5).

Lemma F.4. *Suppose that $\lim_{\varepsilon \rightarrow 0} \mathcal{C}_\varepsilon = \mathcal{C}_0$. Let μ_ε be the optimizer of (εEOT) . Then,*

$$\mu_\varepsilon \rightarrow \mu_* \text{ weakly as } \varepsilon \downarrow 0, \quad \text{and} \quad H(\mu_\varepsilon) \rightarrow H(\mu_*),$$

where $\mu_* = \arg \min_{\mu \in \mathcal{M}_{opt}} H(\mu)$.

Proof. The additive form of (εEOT) and the optimality of the couplings imply that

$$H(\mu_\varepsilon) \leq H(\mu_{\varepsilon'}) \quad \text{and} \quad \int cd\mu_\varepsilon \geq \int cd\mu_{\varepsilon'} \quad \text{for } \varepsilon \geq \varepsilon' > 0.$$

Denote $\mathcal{Q} := \mathcal{M}_{opt}$ and

$$\mathcal{Q}_\varepsilon := \left\{ \mu \in \mathcal{M} : \int cd\mu \leq \int cd\mu_\varepsilon \right\}.$$

Note that \mathcal{Q}_ε is a closed convex set, and $\mu_\varepsilon = \arg \min_{\mu \in \mathcal{Q}_\varepsilon} H(\mu)$.¹⁵ Then $\int cd\mu_\varepsilon \geq \int cd\mu_{\varepsilon'}$ implies that $\mathcal{Q}_\varepsilon \supset \mathcal{Q}_{\varepsilon'}$ for $\varepsilon \geq \varepsilon'$. Next, we claim that $\mathcal{Q} = \bigcap_\varepsilon \mathcal{Q}_\varepsilon$. It is easy to see $\mathcal{Q} \subset \bigcap_\varepsilon \mathcal{Q}_\varepsilon$. For the other direction, take any $\mu \in \bigcap_\varepsilon \mathcal{Q}_\varepsilon$. Then we have $\mu \in \mathcal{Q}$ because

$$\int cd\mu \leq \int cd\mu_\varepsilon \leq \mathcal{C}_\varepsilon \rightarrow \mathcal{C}_0.$$

Then applying Lemma F.3 completes the proof. \square

Thus, it remains to show that $\lim_{\varepsilon \rightarrow 0} \mathcal{C}_\varepsilon = \mathcal{C}_0$. The next result is an adaptation of Nutz (2022, Lemma 5.2).

Lemma F.5. *Suppose that given $\eta > 0$, there exists $\mu^\eta \in \mathcal{M}$ with $\int cd\mu^\eta \leq \mathcal{C}_0 + \eta$ and $H(\mu^\eta) < \infty$. Then $\lim_{\varepsilon \rightarrow 0} \mathcal{C}_\varepsilon = \mathcal{C}_0$.*

Proof. Given $\eta > 0$, we have

$$\mathcal{C}_\varepsilon \leq \int cd\mu^\eta + \varepsilon H(\mu^\eta) \leq \mathcal{C}_0 + \eta + \varepsilon H(\mu^\eta).$$

Thus $\lim_{\varepsilon \rightarrow 0} \mathcal{C}_\varepsilon \leq \mathcal{C}_0 + \eta$. Since $\eta > 0$ is arbitrary, we are done. \square

The next result is an adaptation of Nutz (2022, Lemma 5.4).

Lemma F.6. *Let c be continuous and bounded. Then $\lim_{\varepsilon \rightarrow 0} \mathcal{C}_\varepsilon = \mathcal{C}_0$.*

Proof. Let $\eta > 0$ and $\mu \in \mathcal{M}$ an optimal transport for (OT). By Nutz (2022, Lemma 5.3), there exists $\mu^\eta \in \mathcal{M}$ such that

$$\left| \int cd\mu^\eta - \int cd\mu \right| \leq \eta.$$

Note that $H(\mu^\eta) < \infty$. Then applying Lemma F.5 yields the conclusion. \square

Then by Lemma F.4, the conclusion of Proposition 3.3 follows.

¹⁵Suppose, by contradiction, that there exists $\mu \in \mathcal{Q}_\varepsilon$ such that $H(\mu) < H(\mu_\varepsilon)$. Then

$$\int cd\mu + \varepsilon H(\mu) < \int cd\mu_\varepsilon + \varepsilon H(\mu_\varepsilon),$$

which contradicts with the optimality of μ_ε .

APPENDIX G. OPTIMALITY IN SEMIPARAMETRIC MODELS

We generalize the setup presented in the main texts to allow more flexible sampling distributions for observable data. Our setup here basically follows Christensen *et al.* (2026, Section 5). Assume that data $Z^n = (Z_1, \dots, Z_n)$ are i.i.d. and Z_i follows the distribution $P_{\theta, \eta}$ indexed by $\theta \in \Theta \subset \mathbb{R}^k$ and $\eta \in \mathcal{H}$, where η is a possibly infinite-dimensional nuisance parameter. For instance, in a GMM model, \mathcal{H} is the set of marginal distributions η of Z_i where for each $\eta \in \mathcal{H}$, there exists $\theta \in \Theta$ such that η satisfies the moment restriction $\int g(\theta, z) d\eta = 0$, given some known vector function g .

It is said that $\mathcal{P} = \{P_{\theta, \eta} : \theta \in \Theta, \eta \in \mathcal{H}\}$ has the *least favorable submodels* at (θ, η) if there exist an open neighborhood $\Theta_{\theta, \eta}$ of θ and a map $\Theta_{\theta, \eta} \ni t \mapsto \eta_t \in \mathcal{H}$ such that the parametric submodel $\{P_{t, \eta_t} : t \in \Theta_{\theta, \eta}\}$ has the density function p_{t, η_t} with respect to a common dominating measure ν and satisfies the DQM condition

$$\int \left[\sqrt{p_{\theta+h, \eta_\theta}} - \sqrt{p_{\theta, \eta_\theta}} - \frac{1}{2} h^\top \dot{\ell}_{\theta, \eta} p_{\theta, \eta_\theta} \right]^2 d\nu = o(\|h\|^2), \quad \text{as } h \rightarrow 0,$$

where $\dot{\ell}_{\theta, \eta} : \mathcal{Z}^n \rightarrow \mathbb{R}^k$ is the efficient score function for θ . Thus, the parametric submodel $\{P_{t, \eta_t} : t \in \Theta_{\theta, \eta}\}$ achieves the semiparametric efficiency bound by the inverse of $I_{\theta, \eta} := \int \dot{\ell}_{\theta, \eta} \dot{\ell}_{\theta, \eta}^\top dP_{\theta, \eta_\theta}$. For each (θ, η) , the least favorable submodels need not to be unique. Picking one of them gives no loss of generality because they all behave in the same manner asymptotically.

Following the parametric model, we assume that the planner's utility function w only depends on θ , and not on the nuisance parameter η .

G.1. Decision theoretic framework and rules. Fix $(\theta_0, \eta_0) \in \Theta \times \mathcal{H}$. Consider a least favorable submodel $\{P_{\beta(t)} : t \in \Theta_{\theta_0, \eta_0}\}$, where $\beta(t) = (t, \eta_t)$. Under the reparametrization $t = \theta_0 + h/\sqrt{n} = \theta_{nh}$, Z^n follows the distribution $P_{\beta(\theta_{nh})}^n$. We denote $\overset{h}{\rightsquigarrow}$ by the weak convergence along the path $P_{\beta(\theta_{nh})}^n$, $\overset{h}{\rightarrow}$ by the convergence in probability along $P_{\beta(\theta_{nh})}^n$, and $\overset{0}{\rightarrow}$ by the convergence in probability along P_{θ_0, η_0}^n .

We define the class of the sequences of rules by

$$\mathcal{D} := \left\{ \{\mu_n\} : \mu_n(Z^n) \overset{h}{\rightsquigarrow} Q_{\theta_0, h} \text{ and } \sqrt{n} P_{\beta(\theta_{nh})}^n(\mu_n \notin A_0^\delta) \rightarrow 0 \text{ as } n \rightarrow \infty \quad \forall h \in \mathbb{R}^k, \forall \delta > 0, \forall \theta_0 \in \Theta \right\}.$$

As in the parametric case, a sequence $\{\mu_n\} \in \mathcal{D}$ is said to be *optimal in the limit experiment* if, for each fixed $h \in \mathbb{R}^k$, it converges in distribution along $P_{\beta(\theta_{nh})}^n$ to a coupling μ_∞ that solves $\min_{\mu \in A_0} L_\infty(\mu, \Delta)$ for any realization of the signal $\Delta \sim N(h, I_{\theta_0, \eta_0}^{-1})$.

G.2. Quasi-Bayesian implementation of the Bayesian rules. We replace a posterior function specified in the parametric setup by a quasi-posterior. Let $\hat{\theta}_n$ be the (semiparametrically) efficient estimator of θ , and \hat{I}_n^{-1} be a consistent estimator of the asymptotic covariance $I_{\theta, \eta}^{-1}$. We combine the limited-information quasi-likelihood $N(\hat{\theta}_n, (n\hat{I}_n)^{-1})$ for θ and a prior π on Θ to obtain the quasi-posterior

$$\pi_n(\theta | Z^n) \propto \exp \left(-\frac{1}{2} (\theta - \hat{\theta}_n)^\top (n\hat{I}_n) (\theta - \hat{\theta}_n) \right) \pi(\theta).$$

We compute the Bayesian rules using $\pi_n(\theta|Z^n)$; i.e.,

$$\mu_n^B(z) \in \mathcal{M}_{opt}(z) := \arg \max_{\mu \in \mathcal{M}} \int \sqrt{n}W(\theta, \mu) d\pi_n(\theta|z).$$

Following the parametric model, we construct a unique $\{\mu_n^B(z)\}$ where $\mu_n^B(z)$ minimizes the penalty function H over $\mathcal{M}_{opt}(z)$.

G.3. Optimality results. As an analog for Assumption 3.1 in the parametric models, we impose the following assumptions.

Assumption G.1. (i) Θ is open.

(ii) \mathcal{P} has a least favorable submodel at each $(\theta_0, \eta_0) \in \Theta \times \mathcal{H}$.

(iii) I_{θ_0, η_0} is finite and nonsingular at each $(\theta_0, \eta_0) \in \Theta \times \mathcal{H}$.

(iv) For each $(\theta_0, \eta_0) \in \Theta \times \mathcal{H}$ and each $h \in \mathbb{R}^k$, (iv-a) $\sqrt{n}P_{\beta(\theta_{nh})}^n(\|\hat{\theta}_n - \theta_0\| > \varepsilon) \rightarrow 0$ for each $\varepsilon > 0$ as $n \rightarrow \infty$, and (iv-b) there exists $c \in (0, 1)$ such that $\sqrt{n}P_{\beta(\theta_{nh})}^n(c \leq \hat{\lambda}_{\min}, \hat{\lambda}_{\max} \leq c^{-1}) \rightarrow 0$ as $n \rightarrow \infty$,

(v) For each $(\theta_0, \eta_0) \in \Theta \times \mathcal{H}$, (v-a) $\sqrt{n}(\hat{\theta}_n - \theta_0) \xrightarrow{h} Z$ with $Z \sim N(h, I_{\theta_0, \eta_0}^{-1})$ for all $h \in \mathbb{R}^k$ as $n \rightarrow \infty$, and (v-b) $\hat{I}_n \xrightarrow{0} I_{\theta_0, \eta_0}$ as $n \rightarrow \infty$.

Theorem G.1. Under Assumptions G.1 and 3.2–3.5, $\{\mu_n^B\} \in \mathcal{D}$ is optimal in the limit experiment.

Proof. Once we fix the parameter $(\theta_0, \eta_0) \in \Theta \times \mathcal{H}$, the least favorable submodel $\{P_{\beta(t)} : t \in \Theta_{\theta_0, \eta_0}\}$ becomes a parametric model. Hence, only slight modifications from the proof of Theorem 3.4 are needed. Specifically, Lemmas A.3, A.4, and A.5 follow in the same manner, since their proofs depend only on the properties of w and \mathcal{M} , not on the specific form of the posterior. For Lemma A.2 (i), we need a modification to show $\sqrt{n}P_{\beta(\theta_{nh})}^n(\mu_n^B(Z^n) \notin A_0^\delta) \rightarrow 0$ for all $\delta > 0$ and all $h \in \mathbb{R}^k$, which is given in Lemma G.2 below. For Lemmas A.6 and A.7, we need to use the quasi-posterior counterparts of the Bernstein-von Mises theorem given by Christensen *et al.* (2026, Lemma 5) and the quasi-posterior counterpart of Proposition A.1 given by Xu (2024, Lemma A.5). Auxiliary lemmas given in Appendix B do not need modifications. \square

Lemma G.2. The Bayesian rule $\{\mu_n^B(Z^n)\}$ satisfies $\sqrt{n}P_{\beta(\theta_{nh})}^n(\mu_n^B(Z^n) \notin A_0^\delta) \rightarrow 0$ for all $\delta > 0$ as $n \rightarrow \infty$.

Proof. By the same argument as in the proof of Lemma A.2 (i), for any $\theta_0 \in \Theta$ and any $\delta > 0$, there are \bar{n} and $\varepsilon'_n(\bar{n})$ (which is at the order of $n^{\alpha+1}$ for some $\alpha \geq 1$) such that for all $n \geq \bar{n}$,

$$P_{\beta(\theta_{nh})}^n(\mu_n^B(z) \notin A_0^\delta) \leq P_{\beta(\theta_{nh})}^n(\pi_n(N_{1/n}(\theta_0)^c|z) > 2\varepsilon'_n).$$

Under Assumption G.1 (iii) and (iv), Christensen *et al.* (2026, Lemma 12) implies that

$$\sqrt{n}P_{\beta(\theta_{nh})}^n(\pi_n(N_{1/n}(\theta_0)^c|z) > 2\varepsilon'_n) \rightarrow 0$$

as $n \rightarrow \infty$. \square

REFERENCES

- Adjaho, C. and T. Christensen (2023). “Externally valid policy choice.” arXiv: [2205.05561 \[econ.EM\]](#).
- Angrist, J., E. Bettinger, and M. Kremer (2006). “Long-term educational consequences of secondary school vouchers: Evidence from administrative records in Colombia.” *American Economic Review* 96.3, 847–862. DOI: [10.1257/aer.96.3.847](#).
- Angrist, J., E. Bettinger, and M. Kremer (2019). “Replication data for: Long-term educational consequences of secondary school vouchers: Evidence from administrative records in colombia.” *American Economic Association*. DOI: [10.3886/E116225V1](#).
- Aradillas Fernández, A., J. L. Montiel Olea, C. Qiu, J. Stoye, and S. Tinda (2024). “Robust Bayes treatment choice with partial identification.” arXiv: [2408.11621 \[econ.EM\]](#).
- Athey, S. and S. Wager (2021). “Policy learning with observational data.” *Econometrica* 89.1, 133–161. DOI: [10.3982/ECTA15732](#).
- Banerjee, A. V., S. Chassang, S. Montero, and E. Snowberg (2020). “A theory of experimenters: Robustness, randomization, and balance.” *American Economic Review* 110.4, 1206–1230. DOI: [10.1257/aer.20171634](#).
- Bhattacharya, D. and P. Dupas (2012). “Inferring welfare maximizing treatment assignment under budget constraints.” *Journal of Econometrics* 167.1, 168–196. DOI: [10.1016/j.jeconom.2011.11.007](#).
- Chen, H. and P. Guggenberger (2025). “A note on minimax regret rules with multiple treatments in finite samples.” *Econometric Theory*, 1–27. DOI: [10.1017/S0266466625000040](#).
- Chewi, S., J. Niles-Weed, and P. Rigollet (2024). “Statistical optimal transport.” arXiv: [2407.18163 \[math.ST\]](#).
- Christensen, T., H. R. Moon, and F. Schorfheide (2026). “Optimal decision rules when payoffs are partially identified.” *Review of Economic Studies*. forthcoming. DOI: [10.1093/restud/rdag017](#).
- Clarke, B. S. and A. R. Barron (1990). “Information-theoretic asymptotics of Bayes methods.” *IEEE Transactions on Information Theory* 36.3, 453–471. DOI: [10.1109/18.54897](#).
- Dehejia, R. H. (2005). “Program evaluation as a decision problem.” *Journal of Econometrics* 125.1-2, 141–173. DOI: [10.1016/j.jeconom.2004.04.006](#).
- Fang, Z. and A. Santos (2019). “Inference on directionally differentiable functions.” *Review of Economic Studies* 86.1, 377–412. DOI: [10.1093/restud/rdy049](#).
- Flamary, R., N. Courty, A. Gramfort, M. Z. Alaya, A. Boisbunon, S. Chambon, L. Chapel, A. Corenflos, K. Fatras, N. Fournier, L. Gautheron, N. T. Gayraud, H. Janati, A. Rakotomamonjy, I. Redko, A. Rolet, A. Schutz, V. Seguy, D. J. Sutherland, R. Tavenard, A. Tong, and T. Vayer (2021). “POT: Python optimal transport.” *Journal of Machine Learning Research* 22.78, 1–8.
- Ghosal, S. and A. W. van der Vaart (2017). *Fundamentals of Nonparametric Bayesian Inference*. Vol. 44. Cambridge University Press.
- Ghosh, J. K. and R. V. Ramamoorthi (2003). *Bayesian Nonparametrics*. Springer New York.
- Giacomini, R. and T. Kitagawa (2021). “Robust Bayesian inference for set-identified models.” *Econometrica* 89.4, 1519–1556. DOI: [10.3982/ECTA16773](#).
- Gilboa, I. and D. Schmeidler (1989). “Maxmin expected utility with non-unique prior.” *Journal of Mathematical Economics* 18.2, 141–153. DOI: [10.1016/0304-4068\(89\)90018-9](#).
- Guggenberger, P., N. Mehta, and N. Pavlov (2024). “Minimax regret treatment rules with finite samples when a quantile is the object of interest.” *Working Paper*.
- Hazard, Y. and T. Kitagawa (2025). “Who with whom? Learning optimal matching policies.” arXiv: [2507.13567 \[econ\]](#).
- Heinonen, J. (2001). *Lectures on Analysis on Metric Spaces*. Springer New York, NY.

- Hirano, K. and J. R. Porter (2009). “Asymptotics for statistical treatment rules.” *Econometrica* 77.5, 1683–1701. DOI: [10.3982/ECTA6630](https://doi.org/10.3982/ECTA6630).
- Ishihara, T. and T. Kitagawa (2024). “Evidence aggregation for treatment choice.” arXiv: [2108.06473v2](https://arxiv.org/abs/2108.06473v2) [econ.EM].
- Kido, D. (2022). “Distributionally robust policy learning with wasserstein distance.” arXiv: [2205.04637](https://arxiv.org/abs/2205.04637) [econ].
- Kido, D. (2023). “Locally asymptotically minimax statistical treatment rules under partial identification.” arXiv: [2311.08958](https://arxiv.org/abs/2311.08958) [econ.EM].
- Kim, J.-Y. (2002). “Limited information likelihood and Bayesian analysis.” *Journal of Econometrics* 107.1-2, 175–193. DOI: [10.1016/S0304-4076\(01\)00119-1](https://doi.org/10.1016/S0304-4076(01)00119-1).
- Kitagawa, T., S. Lee, and C. Qiu (2024). “Treatment choice with nonlinear regret.” arXiv: [2205.08586](https://arxiv.org/abs/2205.08586) [econ.EM].
- Kitagawa, T. and A. Tetenov (2018). “Who should be treated? Empirical welfare maximization methods for treatment choice.” *Econometrica* 86.2, 591–616. DOI: [10.3982/ECTA13288](https://doi.org/10.3982/ECTA13288).
- Manski, C. F. (2004). “Statistical treatment rules for heterogeneous populations.” *Econometrica* 72.4, 1221–1246. DOI: [10.1111/j.1468-0262.2004.00530.x](https://doi.org/10.1111/j.1468-0262.2004.00530.x).
- Manski, C. F. (2007). “Minimax-regret treatment choice with missing outcome data.” *Journal of Econometrics* 139.1, 105–115. DOI: [10.1016/j.jeconom.2006.06.006](https://doi.org/10.1016/j.jeconom.2006.06.006).
- Masten, M. A. (2023). “Minimax-regret treatment rules with many treatments.” *Japanese Economic Review* 74.4, 501–537. DOI: [10.1007/s42973-023-00147-0](https://doi.org/10.1007/s42973-023-00147-0).
- Mbakop, E. and M. Tabord-Meehan (2021). “Model selection for treatment choice: Penalized welfare maximization.” *Econometrica* 89.2, 825–848. DOI: [10.3982/ECTA16437](https://doi.org/10.3982/ECTA16437).
- Montiel Olea, J. L., C. Qiu, and J. Stoye (2026). “Decision theory for treatment choice problems with partial identification.” *Review of Economic Studies*. forthcoming. DOI: [10.1093/restud/rdag015](https://doi.org/10.1093/restud/rdag015).
- Nutz, M. (2022). “Introduction to entropic optimal transport.”
- Qi, Z., J.-S. Pang, and Y. Liu (2023). “On robustness of individualized decision rules.” *Journal of the American Statistical Association* 118.543, 2143–2157. DOI: [10.1080/01621459.2022.2038180](https://doi.org/10.1080/01621459.2022.2038180).
- Römisch, W. (2004). “Delta method, infinite dimensional.” *Encyclopedia of Statistical Sciences* 3.
- Shapiro, A. (1990). “On concepts of directional differentiability.” *Journal of Optimization Theory and Applications* 66, 477–487.
- Si, N., F. Zhang, Z. Zhou, and J. Blanchet (2020). “Distributionally robust policy evaluation and learning in offline contextual bandits.” In: *International Conference on Machine Learning*. PMLR, 8884–8894.
- Stoye, J. (2009). “Minimax regret treatment choice with finite samples.” *Journal of Econometrics* 151.1, 70–81. DOI: [10.1016/j.jeconom.2009.02.013](https://doi.org/10.1016/j.jeconom.2009.02.013).
- Stoye, J. (2012). “Minimax regret treatment choice with covariates or with limited validity of experiments.” *Journal of Econometrics* 166.1, 138–156. DOI: [10.1016/j.jeconom.2011.06.012](https://doi.org/10.1016/j.jeconom.2011.06.012).
- Sun, L. (2024). “Empirical welfare maximization with constraints.” arXiv: [2103.15298](https://arxiv.org/abs/2103.15298) [econ.EM].
- Tetenov, A. (2012). “Statistical treatment choice based on asymmetric minimax regret criteria.” *Journal of Econometrics* 166.1, 157–165. DOI: [10.1016/j.jeconom.2011.06.013](https://doi.org/10.1016/j.jeconom.2011.06.013).
- van der Vaart, A. W. (1991). “An asymptotic representation theorem.” *International Statistical Review* 59.1, 97–121. DOI: [10.2307/1403577](https://doi.org/10.2307/1403577).
- van der Vaart, A. W. (1998). *Asymptotic Statistics*. Vol. 3. Cambridge University Press.
- van der Vaart, A. W. and J. A. Wellner (1996). *Weak Convergence and Empirical Processes with Applications to Statistics*. Springer-Verlag. DOI: [10.1007/978-1-4757-2545-2](https://doi.org/10.1007/978-1-4757-2545-2).
- Villani, C. (2009). *Optimal Transport: Old and New*. Vol. 338. Springer.
- Xu, H. (2024). “Asymptotic analysis of point decisions with general loss functions.” *Working Paper*.
- Yata, K. (2023). “Optimal decision rules under partial identification.” arXiv: [2111.04926](https://arxiv.org/abs/2111.04926) [econ.EM].

DEPARTMENT OF ECONOMICS, UNIVERSITY OF ROCHESTER, ROCHESTER, NY 14627, USA.

Email address: ksunada@ur.rochester.edu

DEPARTMENT OF ECONOMICS, UNIVERSITY OF ROCHESTER, ROCHESTER, NY 14627, USA.

Email address: kizumi3@ur.rochester.edu