

Asymptotic Inference for Multi-Stage Stationary Treatment Policy with Variable Selection

Daiqi Gao

DGAO@FAS.HARVARD.EDU

*Department of Statistics
Harvard University
Cambridge, MA 02138, USA*

Yufeng Liu

YFLIU@EMAIL.UNC.EDU

*Department of Statistics and Operations Research, Department of Genetics,
Department of Biostatistics
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599, USA*

Donglin Zeng

DZENG@UMICH.EDU

*Department of Biostatistics
University of Michigan
Ann Arbor, MI 48109, USA*

Editor: Ji Zhu

Abstract

Dynamic treatment regimes or policies are a sequence of decision functions over multiple stages that are tailored to individual features. One important class of treatment policies in practice, namely multi-stage stationary treatment policies, prescribes treatment assignment probabilities using the same decision function across stages, where the decision is based on the same set of features consisting of time-evolving variables (e.g., routinely collected disease biomarkers). Although there has been extensive literature on constructing valid inference for the value function associated with dynamic treatment policies, little work has focused on the policies themselves, especially in the presence of high-dimensional features. We aim to fill the gap in this work. Specifically, we first obtain the multi-stage stationary treatment policy by minimizing the negative augmented inverse probability weighted estimator of the value function to increase asymptotic efficiency. An L_1 penalty is applied on the policy parameters to select important features. We then construct one-step improvements of the policy parameter estimators for valid inference. Theoretically, we show that the improved estimators are asymptotically normal, even if nuisance parameters are estimated at a slow convergence rate and the dimension of the features increases with the sample size. Our numerical studies demonstrate that the proposed method estimates a sparse policy with a near-optimal value function and conducts valid inference for the policy parameters.

Keywords: Augmented inverse probability weighted estimator, dynamic treatment regime, high-dimensional inference, policy parameter, sparse estimation

1. Introduction

Dynamic treatment regimes (DTRs) or policies have recently attracted considerable attention in precision medicine. They are a sequence of decision functions to prescribe treatments

over stages based on an individual’s features which can evolve over time. When the individual’s features consist of clinical variables that are routinely collected over stages, one important class of policies, which we name as multi-stage stationary treatment policies (MSTPs), are to prescribe from the same set of treatments using the same decision function over all stages. Moreover, the policies are stochastic, meaning that individuals receive recommended treatments with a probability close but not equal to one. An MSTP is particularly critical for delivering interventions in mobile health (mHealth), where users receive treatments across multiple treatment stages, and the policy’s effectiveness over all stages must be optimized. For example, OhioT1DM (Marling and Bunescu, 2020) is a dataset developed to support health research for individuals with Type 1 Diabetes (T1D). The physiological and behavioral information, including the current blood glucose level, meal times, exercise intensity, and other self-reported events, is available at each stage. An MSTP that considers these features can guide decisions on insulin delivery, helping to maintain blood glucose levels within the target range over the long term. We focus on MSTPs in this paper because they are particularly useful for treating chronic diseases in clinical practices. First, the policies use the same decision function and the same set of variables so they are convenient for both implementation and interpretation. Second, the policies are dynamic by incorporating individual’s evolving features in decisions, which are often known to be important for disease prognostics and thus are routinely collected during clinical visits or lab tests. Third, using stochastic decisions or stochastic policies is more flexible than relying on a deterministic policy whose estimation is known to be sensitive to evidence bias and noise. Moreover, stochastic policies can be deployed in future studies, as they satisfy the positivity assumption required for estimating treatment effects (Boruvka et al., 2018) and support efficient offline policy learning (Gao et al., 2022) in between-study or after-study analysis.

Many approaches have been developed to estimate optimal DTRs using data from a multi-stage study. For example, Q-learning derives the DTR by maximizing the optimal Q-function, which is the expected optimal cumulative reward for each treatment given the current features (Murphy, 2005; Zhao et al., 2011; Moodie et al., 2012; Zhu et al., 2019). The Q-function is typically estimated using regression in a backward fashion when the number of treatment stages (also called the horizon) is finite. On the other hand, A-learning poses assumptions only on the interaction effect between the treatments and actions, and thus avoids potential problems of misspecified treatment-free effect (Murphy, 2003; Shi et al., 2018; Jeng et al., 2018). Policy search methods find the DTR within a function class that directly maximizes the value function, which is estimated using inverse probability weighted (IPW) estimator (Zhao et al., 2015) or the augmented inverse propensity weighted estimator (AIPWE) (Zhang et al., 2013; Liu et al., 2018b). The latter is often called the doubly robust estimator. However, these methods typically focus on a small number of treatment stages where the environment is generally non-stationary and a time-dependent policy is preferred. In the reinforcement learning (RL) literature, Nie et al. (2021) proposed learning a non-stationary policy using a doubly robust estimator for finite-horizon problems without assuming Markovian dynamics, and provide bounds on the convergence rate of the advantage function. Value iteration algorithms (Antos et al., 2007) and doubly robust methods (Luckett et al., 2020; Liao et al., 2022) have been proposed to find the optimal stationary policy in infinite horizon Markov decision processes (MDPs) using off-policy data.

A comprehensive review of DTRs and RL can be found in Kosorok and Laber (2019) and Sutton and Barto (2018). However, none of the above provide inferential results for the policy.

In addition to DTR estimation, a number of studies have considered obtaining valid inference for the value function associated with the estimated DTR. For example, Luedtke and Van Der Laan (2016); Zhu et al. (2019) and Shi et al. (2020a) studied value inference, allowing for situations where the treatment is neither beneficial nor harmful for a subpopulation. Similar inference has been investigated in the RL framework under MDP assumptions (Kallus and Uehara, 2020; Lueckett et al., 2020; Liao et al., 2021; Shi et al., 2021, 2022; Kallus and Uehara, 2022). There are few methods to conduct inference for the treatment policies themselves. This is especially important for studying MSTPs with many features: clinicians typically favor simple and parsimonious decisions due to concerns about implementation and the cost of collecting disease biomarkers during routine visits or lab tests. Jeng et al. (2018); Zhu et al. (2019) derived the asymptotic distribution of the parameters in Q-learning or A-learning under the assumption that the regression models are correctly specified, which may not be plausible in the presence of many features. More recently, Liang et al. (2022) studied inference for deterministic treatment decisions in a high-dimensional setting but restricted their analysis to a single stage.

This work aims to address the gap by obtaining valid inference for MSTPs with high-dimensional features. We assume the data come from micro-randomized trials (MRTs), where treatments are randomly assigned according to a known behavior policy at each stage. Our stationary treatment policy is defined as a probability function of a linear combination of the features, and a tuning parameter is used in this function to approximate a deterministic decision. For inference, we first estimate the value function using the AIPWE (Bang and Robins, 2005), with a discussion on different constructions of the augmentation term. Jiang and Li (2016); Thomas and Brunskill (2016) first proposed this estimator for off-policy evaluation (OPE) of a fixed target policy. Kallus and Uehara (2020) further proved that this estimator based on the ratio between the behavior policy and the target policy achieves semiparametric efficiency in non-Markovian decision processes, and proposed an improved estimator that uses a marginalized density ratio of the current state and action to achieve semiparametric efficiency for MDPs. Furthermore, to find a sparse estimator of MSTP, we impose an L_1 penalty for variable selection. The classical inferential theory fails in this case due to the presence of high-dimensional parameters. The asymptotic distribution of the estimated parameter becomes intractable due to the non-ignorable estimation bias and the sparsity effect of the nuisance parameters. To validate the inference for the parameter estimators, we adopt the idea of one-step estimation (Zhang and Zhang, 2014; Ning and Liu, 2017) to remove bias in the regularized estimators.

Our main contributions are summarized as follows. (1) We learn a sparse policy based on high-dimensional features and conduct statistical inference on it. To the best of our knowledge, this is the first work that considers inference for multi-stage treatment rules themselves within policy search methods, especially in high-dimensional settings. (2) Theoretically, we show that the final estimators for the policy parameters are asymptotically normal even if the dimension of the features increases exponentially with \sqrt{n} , or the models for Q -functions are misspecified and their parameters are estimated at a rate arbitrarily slow. The theoretical analysis is more involved than general high-dimensional inference problems as in

Ning and Liu (2017), since we need to account for the high-dimensional plug-in estimator of the nuisance parameters and the high-dimensional policy parameters. The augmentation in our first step shares a similar spirit with debiased machine learning methods in Chernozhukov et al. (2018a), where nuisance parameters (Q-functions) are decorrelated from the policy parameters, except that we do not rely on data-splitting to ensure stable estimation with limited data. The one-step improvement in our second step further decorrelates the parameter of interest from the remaining high-dimensional parameters in MSTP. In other words, our proposed method incorporates two decorrelation procedures for inference. (3) The optimization and inference of multi-stage policies with off-policy data are computationally challenging. The importance sampling weight (Liu and Lee, 2017), also known as the cumulative density ratio (Kallus and Uehara, 2020), is used to adjust for discrepancies between the behavior policy used to collect the data and the target policy being optimized. The variance of the weights increases with the number of treatment stages. We leverage several techniques to stabilize the estimated asymptotic variance, including using the difference quotient (Lax and Terrell, 2014) to estimate the gradient and Hessian matrix of the loss function.

The rest of this paper is organized as follows. In Section 2, we define the objective of our problem and describe our method for learning the MSTP. Then we introduce the procedure for constructing the one-step estimator in theory and specify some implementation details. In Section 3, we show that the one-step estimator is asymptotically normal and provide its confidence intervals (CI). In Sections 4 and 5, we demonstrate our method in simulation studies and a real data example. Finally, we conclude this paper with some discussion in Section 6.

2. Methodology

We consider selecting treatment policies for T stages, where the feature space \mathcal{X} and the treatment space $\mathcal{A} = \{-1, 1\}$ are the same across stages. At each stage, $\mathbf{X} = [X_{1:d}] \in \mathcal{X}$ is a d -dimensional feature vector, $A \in \mathcal{A}$ is the treatment, and $R \in \mathbb{R}$ is the reward. Here, we use the notation $i : j$ to represent the index set i, \dots, j , where the indices $i, j \in \mathbb{N}$. The distributions of the feature \mathbf{X} and the treatment A may vary across stages. An MSTP is defined as a mapping $\pi : \mathcal{X} \mapsto \mathcal{P}(\mathcal{A})$ from the feature space to the space of probability distributions over the action space, which assigns treatments with the same function in all stages. Specifically, $\pi(a|\mathbf{X})$ denotes the probability of assigning treatment $a \in \mathcal{A}$ when the feature is \mathbf{X} .

We assume that data are obtained from an offline, T -stage MRT, so the observed trajectory of the i -th subject is denoted as

$$\mathbf{D}_i = \{\mathbf{X}_{i,1}, A_{i,1}, R_{i,1}, \mathbf{X}_{i,2}, A_{i,2}, R_{i,2}, \dots, \mathbf{X}_{i,T}, A_{i,T}, R_{i,T}\},$$

where the reward $R_{i,t} \in \mathbb{R}$ is an unknown function of the data $\{\mathbf{X}_{i,1:t}, A_{i,1:t}, R_{i,1:(t-1)}\}$ observed prior to or at time t . Let the domain of \mathbf{D}_i be defined as $\mathcal{D} := (\mathcal{X} \times \mathcal{A} \times \mathbb{R})^T$. Assume there are n subjects and their trajectories $\{\mathbf{D}_i\}_{i=1}^n$ are independent and identically distributed. In addition, each action $A_{i,t}$ is taken randomly with a known probability $\mu_t(A_{i,t}|\mathbf{H}_{i,t})$, which depends on the history $\mathbf{H}_{i,t} = \{\mathbf{X}_{i,1:t}, A_{i,1:(t-1)}, R_{i,1:(t-1)}\}$. The collection $\boldsymbol{\mu} := [\mu_{1:T}]$ is sometimes called the behavior policy in the RL literature.

An important metric to evaluate an MSTP $\pi^\theta(a|\mathbf{X})$ with parameter θ is called the value function, which is defined as the sum of rewards $V(\theta) = \mathbb{E}^\theta(\sum_{t=1}^T R_t)$. The optimal MSTP in a policy class Π , whose parameter is denoted as θ^* , is the one maximizing $V(\theta)$, i.e., the optimal parameter $\theta^* \in \arg \max_{\theta: \pi^\theta \in \Pi} V(\theta)$. We focus on the policy class

$$\Pi = \left\{ \pi^\theta : \pi^\theta(a|\mathbf{X}) = \frac{e^{ag(\mathbf{X}, \theta)/\tau}}{1 + e^{ag(\mathbf{X}, \theta)/\tau}}, \|\theta\|_2 \leq 1, a \in \mathcal{A}, \mathbf{X} \in \mathcal{X} \right\},$$

where $\theta = [\theta_{0:d}]$ and τ is a constant scaling parameter. The function g is taken to be linear, defined as $g(\mathbf{X}, \theta) = \theta_0 + \sum_{j=1}^d X_j \theta_j$. It provides direct interpretability regarding the influence of each feature and ensures estimation efficiency with limited sample sizes, both of which are essential in clinical practices. Note that the scaling parameter τ adjusts the strength of influence of parameters on the action probability. A smaller τ leads to a decision closer to a deterministic policy. The constraint $\|\theta\|_2 = 1$ ensures boundedness of the optimal policy within the class Π . Without such constraints, the optimal policy π^{θ^*} will have $\theta^* \rightarrow \infty$ to approximate the deterministic policy. Without confusion, we use \mathbb{E}^θ to denote the expectation under the policy indexed by θ , and use \mathbb{E} to denote the expectation under the behavior policy. Our goal is to estimate the optimal MSTP and obtain a proper inference for the policy parameter θ .

To infer the optimal MSTP using the observed data, we make the following three standard assumptions (Robins, 2000; Murphy, 2003; Nie et al., 2021). Denote $\mathbf{X}_{t+1}(a_{1:t})$ and $R_t(a_{1:t})$ as the potential outcomes if an action sequence $a_{1:t} \in \mathcal{A}^t$ were taken.

Assumption 1 (Sequential ignorability) *The potential outcomes $\{\mathbf{X}_{t'+1}(A_{1:(t-1)}, a_{t:t'}), R_{t'}(A_{1:(t-1)}, a_{t:t'})\}_{t'=t}^T$ is independent of the treatment A_t given $\{\mathbf{X}_{1:t}, A_{1:(t-1)}\}$ for all $a_t \in \mathcal{A}$, $t \in \{1 : T\}$.*

Assumption 2 (Consistency) *The observed outcomes are consistent with the potential outcomes, $\mathbf{X}_{t+1} = \mathbf{X}_{t+1}(A_{1:t})$, $R_t = R_t(A_{1:t})$ for all $t \in \{1 : T\}$.*

Assumption 3 (Positivity) *There exists a constant $p_0 > 0$ such that $\mu_t(a|\mathbf{H}_t) \geq p_0$ for all $a \in \mathcal{A}$ and $t \in \{1 : T\}$.*

Note that Assumption 3 holds if the data are from an MRT, which we assume for the subsequent development.

2.1 Estimate Policy Parameter with Variable Selection

To estimate the value function $V(\theta)$ using the data with known μ_t , we first denote $Q_t(\mathbf{x}_t, a_t)$ as some function of rewards given $\mathbf{X}_t = \mathbf{x}_t$ and $A_t = a_t$. For example, $Q_t(\mathbf{x}_t, a_t) := \mathbb{E}^\theta\{\sum_{k=t}^T R_k \mid \mathbf{X}_t = \mathbf{x}_t, A_t = a_t\}$ represents the Q-function at stage t under the policy π^θ . In general, Q_t may depend on the parameter θ , as in the Q-function of the policy π^θ . In such cases, it is often first obtained based on an initial estimate $\hat{\theta}$, and then plugged in during the subsequent estimation of θ , as will be further discussed in Section 2.3.

Let the expectation of Q_t given \mathbf{X}_t be

$$U_t(\mathbf{x}_t) := \mathbb{E}\{Q_t(\mathbf{X}_t, A_t) \mid \mathbf{X}_t = \mathbf{x}_t\},$$

where $A_t \sim \pi^\theta(\cdot | \mathbf{X}_t)$. Denote $\mathbf{Q} := [Q_{1:T}]$ and $\mathbf{U} := [U_{1:T}]$. Then, based on the AIPWE given in Kallus and Uehara (2020), we construct the following estimator for $V(\theta)$:

$$\hat{V}(\theta) = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \left\{ \rho_{i,1:t}^{\theta,\mu}(\mathbf{H}_{i,t}, A_{i,t}) [R_{i,t} - Q_t(\mathbf{X}_{i,t}, A_{i,t})] + \rho_{i,1:(t-1)}^{\theta,\mu}(\mathbf{H}_{i,t-1}, A_{i,t-1}) U_t(\mathbf{X}_{i,t}) \right\},$$

where $\rho_{i,1:t}^{\theta,\mu}(\mathbf{H}_{i,t}, A_{i,t}) := \prod_{k=1}^t \frac{\pi^\theta(A_{i,k} | \mathbf{X}_{i,k})}{\mu_k(A_{i,k} | \mathbf{H}_{i,k})}$ is often called the importance sampling weight. For each stage t , to adjust for the difference between the behavior policy μ and the target policy π^θ , $\hat{V}(\theta)$ uses the step-wise weight for each reward within a trajectory. Compared to the trajectory-wise weight, which is the weight for the entire trajectory, the step-wise weight marginalizes over the weights for future actions and thus reduces the variance of the value estimator (Jiang and Li, 2016). One major advantage of using the AIPWE for estimating $V(\theta)$ is that the estimator $\hat{V}(\theta)$ is unbiased for arbitrary \mathbf{Q} , as long as U_t is correctly evaluated as the conditional expectation of Q_t . To see this, notice that the expectation of $\hat{V}(\theta)$ is

$$\begin{aligned} \mathbb{E} \hat{V}(\theta) &= \mathbb{E}^\theta \left\{ \sum_{t=1}^T R_{i,t} \right\} - \mathbb{E} \left\{ \sum_{t=1}^T \rho_{i,1:(t-1)}^{\theta,\mu}(\mathbf{H}_{i,t-1}, A_{i,t-1}) \cdot \right. \\ &\quad \left. \mathbb{E} \left[\frac{\pi^\theta(A_{i,t} | \mathbf{X}_{i,t})}{\mu_t(A_{i,t} | \mathbf{H}_{i,t})} Q_t(\mathbf{X}_{i,t}, A_{i,t}) - U_t(\mathbf{X}_{i,t}) \middle| \mathbf{H}_{i,t} \right] \right\} \end{aligned}$$

given any Q_t and U_t . The first term on the right-hand side follows from inverse probability weighting, and the second term follows from the tower property of conditional expectation. Now by the definition of U_t ,

$$\begin{aligned} &\mathbb{E} \left\{ \frac{\pi^\theta(A_{i,t} | \mathbf{X}_{i,t})}{\mu_t(A_{i,t} | \mathbf{H}_{i,t})} Q_t(\mathbf{X}_{i,t}, A_{i,t}) - U_t(\mathbf{X}_{i,t}) \middle| \mathbf{H}_{i,t} \right\} \\ &= \sum_{a \in \mathcal{A}} \mu_t(a | \mathbf{H}_{i,t}) \frac{\pi^\theta(a | \mathbf{X}_{i,t})}{\mu_t(a | \mathbf{H}_{i,t})} Q_t(\mathbf{X}_{i,t}, a) - U_t(\mathbf{X}_{i,t}) = 0 \end{aligned}$$

once $U_t(\mathbf{X}_{i,t}) = \mathbb{E}[Q_t(\mathbf{X}_{i,t}, A_{i,t}) | \mathbf{X}_{i,t}]$. Therefore, $\mathbb{E} \hat{V}(\theta) = V(\theta)$ for any Q_t . In OPE for any fixed policy π in non-Markov decision processes, the value function estimator $\hat{V}(\theta)$ is semiparametrically efficient when Q_t is defined as the Q-function Q_t^θ of π^θ at stage t , and is estimated using sample splitting with a rate condition or without sample splitting with a Donsker condition. Besides, $\hat{V}(\theta)$ is a doubly robust estimator, which guarantees \sqrt{n} -consistency even if one of the nuisance parameters μ and \mathbf{Q} is misspecified in an observational study (Kallus and Uehara, 2020).

The above property implies that we can always obtain a consistent estimator for θ^* with any working model for Q_t . However, a good choice of Q_t may improve the asymptotic efficiency of the estimated policy parameter $\hat{\theta}$. We will discuss the different methods for estimating \hat{Q}_t in Section 2.3. To ensure the relationship between Q_t and U_t to be correct, we estimate U_t using $\hat{U}_t(\mathbf{x}) = \sum_{a \in \mathcal{A}} \pi^\theta(a | \mathbf{x}) \hat{Q}_t(\mathbf{x}, a)$. After plugging \hat{Q}_t and \hat{U}_t into the expression of $\hat{V}(\theta)$, we aim to maximize $\hat{V}(\theta)$, or equivalently, minimize the function

$$\ell(\theta, \hat{\eta}) := \frac{1}{n} \sum_{i=1}^n l_i(\theta, \hat{\eta}), \quad (1)$$

where $\boldsymbol{\eta} := \boldsymbol{Q}$ is the nuisance parameter, and

$$l_i(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}) := - \sum_{t=1}^T \left\{ \rho_{i,1:t}^{\boldsymbol{\theta}, \boldsymbol{\mu}}(\boldsymbol{H}_{i,t}, A_{i,t}) [R_{i,t} - \hat{Q}_t(\boldsymbol{X}_{i,t}, A_{i,t})] + \rho_{i,1:(t-1)}^{\boldsymbol{\theta}, \boldsymbol{\mu}}(\boldsymbol{H}_{i,t-1}, A_{i,t-1}) \hat{U}_t(\boldsymbol{X}_{i,t}) \right\}. \quad (2)$$

Now we have $\mathbb{E}[\ell(\boldsymbol{\theta}, \boldsymbol{\eta})] = -V(\boldsymbol{\theta})$ for any $\boldsymbol{\eta}$. Assume $\boldsymbol{\eta} \in \mathcal{H}$, where \mathcal{H} is a convex subset of some normed vector space. Suppose the estimator $\hat{\boldsymbol{\eta}}$ converges in probability uniformly to some deterministic limit $\bar{\boldsymbol{\eta}}$. Notably, only the limit $\bar{\boldsymbol{\eta}}$, and not the true value $\boldsymbol{\eta}^*$, will appear in the asymptotic distribution of $\hat{\boldsymbol{\theta}}$. Since $\boldsymbol{\theta}$ is high-dimensional, we include an L_1 -penalty to obtain a sparse MSTP. Specifically, we estimate $\boldsymbol{\theta}$ as

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} [\ell(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}) + \lambda_{\boldsymbol{\theta}} \|\boldsymbol{\theta}\|_1] \quad \text{s.t.} \quad \|\boldsymbol{\theta}\|_2 \leq 1, \quad (3)$$

where $\lambda_{\boldsymbol{\theta}}$ is a tuning parameter for $\boldsymbol{\theta}$. The L_1 penalty is primarily chosen due to its convexity and the availability of efficient optimization methods such as proximal gradient descent (Sra et al., 2011). Other nonconvex penalties like MCP (Zhang, 2010) and SCAD (Fan and Li, 2001) can also be incorporated into the proposed method (Ning and Liu, 2017).

2.2 Statistical Inference for Sparse High Dimensional Parameters

To obtain valid inference, we follow Ning and Liu (2017) to construct a one-step improvement of the estimator obtained in the previous section. Note that the classical Rao's score test is based on the asymptotically normal distribution of the profile score function under the null hypothesis. However, in high-dimensional settings, the asymptotic distribution becomes intractable due to the estimation bias in regularized estimators.

When the parameter of interest is θ_j , we denote $\boldsymbol{\nu}_j := (\theta_{1:(j-1)}, \theta_{(j+1):d})$ and let $\boldsymbol{\xi} := (\theta_j, \boldsymbol{\nu}_j)$. Now both $\boldsymbol{\eta}$ and $\boldsymbol{\nu}_j$ can be viewed as nuisance parameters. Let the corresponding optimal parameters of θ_j and $\boldsymbol{\nu}_j$ be θ_j^* and $\boldsymbol{\nu}_j^*$. Given the estimated nuisance parameter $\hat{\boldsymbol{\eta}}$, we can define the matrix $\mathbf{I} := \mathbb{E}[\nabla_{\boldsymbol{\xi}} \boldsymbol{\xi} \ell(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})]$. With $I_{\theta_j \theta_j}, \mathbf{I}_{\theta_j \boldsymbol{\nu}_j}, \mathbf{I}_{\boldsymbol{\nu}_j \boldsymbol{\nu}_j}, \mathbf{I}_{\boldsymbol{\nu}_j \theta_j}$ being the corresponding submatrices of \mathbf{I} with respect to the partition of the parameters, let $I_{\theta_j | \boldsymbol{\nu}_j} = I_{\theta_j \theta_j} - \mathbf{I}_{\theta_j \boldsymbol{\nu}_j} \mathbf{I}_{\boldsymbol{\nu}_j \boldsymbol{\nu}_j}^{-1} \mathbf{I}_{\boldsymbol{\nu}_j \theta_j}$. Next, a decorrelated score function is defined as

$$S_j(\theta_j, \boldsymbol{\nu}_j, \boldsymbol{\eta}) = \nabla_{\theta_j} \ell(\theta_j, \boldsymbol{\nu}_j, \boldsymbol{\eta}) - \boldsymbol{w}_j^T \nabla_{\boldsymbol{\nu}_j} \ell(\theta_j, \boldsymbol{\nu}_j, \boldsymbol{\eta}),$$

where $\boldsymbol{w}_j^T = \mathbf{I}_{\theta_j \boldsymbol{\nu}_j} \mathbf{I}_{\boldsymbol{\nu}_j \boldsymbol{\nu}_j}^{-1}$.

The definition comes from the Taylor expansion of the score function, and ensures that S_j is uncorrelated with the nuisance score function. The decorrelated score function is guaranteed to have asymptotic normality under the null hypothesis that $H_0 : \theta_j^* = 0$, and can be used for hypothesis testing. Given the estimated parameters $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\eta}}$, the sparse estimator $\hat{\boldsymbol{w}}_j$ of \boldsymbol{w}_j can be obtained using the Dantzig estimator

$$\hat{\boldsymbol{w}}_j = \arg \min \|\boldsymbol{w}_j\|_1, \quad \text{s.t.} \quad \left\| \nabla_{\theta_j \boldsymbol{\nu}_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \boldsymbol{w}_j^T \nabla_{\boldsymbol{\nu}_j \boldsymbol{\nu}_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \right\|_{\infty} \leq \lambda_{\boldsymbol{w}_j}, \quad (4)$$

where $\lambda_{\boldsymbol{w}_j}$ is tuned by cross-validation for each j separately. This Dantzig estimator (4 is essentially the best sparse linear combination of the nuisance score function $\nabla_{\boldsymbol{\nu}_j} \ell(\theta_j, \boldsymbol{\nu}_j, \boldsymbol{\eta})$

that approximates the score function $\nabla_{\theta_j} \ell(\theta_j, \nu_j, \eta)$ of the parameter of interest with error $\lambda_{\mathbf{w}_j}$. It can also be estimated by penalized least squares using the gradient or Hessian matrix of ℓ (see Ning and Liu, 2017, Remark 3). Let $\hat{\nu}_j := (1, -\hat{\mathbf{w}}_j^T)^T$. Then $S_j(\theta_j, \nu_j, \hat{\eta})$ can be estimated as

$$\hat{S}_j(\hat{\theta}_j, \hat{\nu}_j, \hat{\eta}) = \nabla_{\theta_j} \ell(\hat{\theta}_j, \hat{\nu}_j, \hat{\eta}) - \hat{\mathbf{w}}_j^T \nabla_{\nu_j} \ell(\hat{\theta}_j, \hat{\nu}_j, \hat{\eta}) \quad (5)$$

by plugging in $\hat{\theta}$, $\hat{\eta}$ and $\hat{\mathbf{w}}_j$.

Finally, building on the sparse estimator $\hat{\theta}$ introduced in the previous section, a one-step estimator is defined as

$$\tilde{\theta}_j := \hat{\theta}_j - \hat{S}_j(\hat{\theta}_j, \hat{\nu}_j, \hat{\eta}) / \hat{I}_{\theta_j|\nu_j}, \quad \text{where } \hat{I}_{\theta_j|\nu_j} = \nabla_{\theta_j}^2 \ell(\hat{\theta}, \hat{\eta}) - \hat{\mathbf{w}}_j^T \nabla_{\nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \quad (6)$$

for $j \in \{1 : d\}$. This estimator solves the equation $\hat{S}_j(\theta_j, \hat{\nu}_j, \hat{\eta}) = 0$ through a computationally efficient one-step method. Directly solving for the root of the decorrelated score function can be computationally demanding, especially when the score function is non-convex or has multiple roots. The one-step method provides an efficient alternative by solving the estimating equation $\hat{S}_j(\hat{\theta}_j, \hat{\nu}_j, \hat{\eta}) + \hat{I}_{\theta_j|\nu_j}(\theta_j - \hat{\theta}_j) = 0$. The resulting estimator is asymptotically equivalent to fully iterative solutions of the decorrelated score function.

For the true parameter θ^* , define the matrix $\mathbf{I}^* := \mathbb{E}[\nabla_{\xi\xi} l_0(\theta^*, \bar{\eta})]$ and the vectors $\mathbf{w}_j^* = \mathbf{I}_{\nu_j\nu_j}^{*-1} \mathbf{I}_{\nu_j\theta_j}^*$, $\mathbf{v}_j^* = (1, -\mathbf{w}_j^{*T})^T$ where l_0 is an independent copy of l_i for any i . Now let

$$\Sigma^* := \text{Var}[\nabla_{\xi} l_0(\theta^*, \bar{\eta})] \quad \text{and} \quad \sigma_j^* := \mathbf{v}_j^{*T} \Sigma^* \mathbf{v}_j^*. \quad (7)$$

In general, Σ^* could be different from \mathbf{I}^* since our loss function is not the negative log-likelihood. As will be shown in Section 3, $\tilde{\theta}_j$ is asymptotically normal with mean θ_j^* and asymptotic variance $\sigma_j^* / [n I_{\theta_j|\nu_j}^{*2}]$, which can be used for constructing the CIs. Note that after the decorrelation, the one-step estimator $\tilde{\theta}_j$ is still a consistent estimator of θ_j^* .

We summarize the theoretical steps of estimating the high-dimensional parameters and making inference with one-step estimators in Algorithm 1. Note that the confidence level applies to each individual CI rather than to multiple CIs simultaneously. Implementation methods will be specified in Section 2.3.

2.3 Implementation

The implementation of the theoretical method described in Algorithm 1 is particularly challenging due to several reasons. First, the importance sampling weight $\rho_{i,1:t}^{\theta, \mu}$ in (1) is unstable in numerical computation since it is a product of t probability ratios. The variance of the value estimator can thus grow exponentially with T (Liu et al., 2018a). This problem is exacerbated in our case because the true optimal parameter θ^* may be away from the behavior policy, further increasing the variance of the loss function. Second, an ideal nuisance parameter \mathbf{Q} might depend on the parameter θ to be optimized. For example, as mentioned in Section 2.1, estimating \mathbf{Q} as the Q-functions of π^θ ensures that the value function estimator $\hat{V}(\theta)$ is semiparametrically efficient (Kallus and Uehara, 2020). However, solving (3) becomes computationally expensive when θ appears both in the policy π^θ and in the Q-function Q_t^θ , and this can further exacerbate numerical instability. Third,

Algorithm 1: Optimization and inference of high-dimensional MSTP parameters**Input** : n trajectories $\mathbf{D}_{1:n}$.**Output:** The sparse estimator $\hat{\boldsymbol{\theta}}$, the one-step estimator $\tilde{\boldsymbol{\theta}}_j$ with its CI for $j \in \{1 : d\}$.

- 1 Estimate $\hat{\mathbf{Q}}$ using an arbitrary working model;
- 2 Estimate the policy parameter $\hat{\boldsymbol{\theta}}$ by solving (3) with $\hat{\mathbf{Q}}$, where $\lambda_{\boldsymbol{\theta}}$ is tuned via cross-validation;
- 3 **for** $j \in \{1 : d\}$ **do**
- 4 Partition the sparse estimate $\hat{\boldsymbol{\theta}}$ as $(\hat{\theta}_0, \hat{\theta}_j, \hat{\boldsymbol{\nu}}_j)$;
- 5 Obtain the Dantzig-type estimator $\hat{\boldsymbol{w}}_j$ by solving (4);
- 6 Compute the decorrelated score function $\hat{S}_j(\theta_j, \hat{\boldsymbol{\nu}}_j, \hat{\boldsymbol{\eta}})$ using (5);
- 7 Calculate the one-step estimator $\tilde{\theta}_j$ using (6);
- 8 Construct the $(1 - \alpha) \times 100\%$ CI for θ_j as

$$\left(\tilde{\theta}_j - \Phi^{-1}(1 - \alpha/2) \frac{\sqrt{\hat{\sigma}_j}}{\sqrt{n} \hat{I}_{\theta_j|\boldsymbol{\nu}_j}}, \tilde{\theta}_j + \Phi^{-1}(1 - \alpha/2) \frac{\sqrt{\hat{\sigma}_j}}{\sqrt{n} \hat{I}_{\theta_j|\boldsymbol{\nu}_j}} \right), \quad (8)$$

where Φ is the cumulative distribution function of the standard normal distribution, and

$$\hat{\sigma}_j = (1, -\hat{\boldsymbol{w}}_j^T) \left[\frac{1}{n} \sum_{i=1}^n \nabla_{\boldsymbol{\xi}} l_i(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \nabla_{\boldsymbol{\xi}} l_i(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}})^T \right] (1, -\hat{\boldsymbol{w}}_j^T)^T.$$

9 **end**

the function (3) itself is difficult to maximize due to the non-convexity of the loss function $\ell(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})$, the non-differentiability of the L_1 penalty, and the L_2 constraint. Examples of non-convex loss functions are provided in Section 4 and Appendix B.4. In this section, we will address these challenges.

Stabilizing the importance sampling weight. The importance sampling weight $\rho_{i,1:t}^{\boldsymbol{\theta}, \boldsymbol{\mu}}$ appears in the loss function $\ell(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})$ in (3) when finding the sparse estimator $\hat{\boldsymbol{\theta}}$, and also appears in the gradient and Hessian matrix of $\ell(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})$ in (4), (5), (6), and (8) when constructing the CIs for $\tilde{\theta}_j$ (see Appendix D.1 for the gradient and Hessian matrix of $\ell(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})$). A practical solution for stabilizing the weight is to use the weighted importance sampling (Thomas and Brunskill, 2016), which scales the probability ratio by the average ratio across all episodes at this stage. Specifically, the sparse estimator $\hat{\boldsymbol{\theta}}$ can be found by

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \left\{ -\frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \frac{\rho_{i,1:t}^{\boldsymbol{\theta}, \boldsymbol{\mu}}}{\omega_{1:t}^{\boldsymbol{\theta}, \boldsymbol{\mu}}} [R_{i,t} - \hat{Q}_t(\mathbf{X}_{i,t}, A_{i,t})] + \frac{\rho_{i,1:(t-1)}^{\boldsymbol{\theta}, \boldsymbol{\mu}}}{\omega_{1:t-1}^{\boldsymbol{\theta}, \boldsymbol{\mu}}} \hat{U}_t(\mathbf{X}_{i,t}) + \lambda_{\boldsymbol{\theta}} \|\boldsymbol{\theta}\|_1 \right\} \quad \text{s.t.} \quad \|\boldsymbol{\theta}\|_2 \leq 1, \quad (9)$$

where $\omega_{1:t}^{\theta, \mu} = \frac{1}{n} \sum_{j=1}^n \rho_{j,1:t}^{\theta, \mu}$ and $\lambda_{\theta}, \lambda_{\theta}$ are tuning parameters for the initial and final sparse estimators, respectively. Similar weighted ratios have been used and discussed in Precup (2000); Thomas (2015); Thomas and Brunskill (2016).

In addition, our simulation shows that the gradient and the Hessian matrix can be more unstable than the loss function due to the product of more terms of probabilities. Therefore, we propose two amendments to solve this problem in practice. First, for the gradient and the Hessian matrix in (4), (5), and (6), we calculate them numerically using the symmetric difference quotient (see Appendix A.4). Since the ℓ_2 -norm of $\hat{\theta}$ is restricted to be 1, we carefully handle the special case when $\hat{\theta}_0 = 0$, where only Newton's difference quotient is available. Second, we use bootstrap to numerically find the CI of $\hat{\theta}_j$ in step (8). That is, we randomly select B bootstrap samples and compute the sparse estimator $\hat{\theta}_b$ and the one-step estimator $\tilde{\theta}_{b,j}$ for $j \in \{1 : d\}$ on each bootstrap sample. The CI of $\tilde{\theta}_j$ is then constructed as the $\alpha/2$ and $1 - \alpha/2$ quantiles of $\{\tilde{\theta}_{b,j}\}_{b=1}^B$, corresponding to a confidence level of $(1 - \alpha)$.

Estimating the nuisance parameters. When the nuisance parameter Q depends on the policy parameter θ , we propose to estimate Q using an initial estimator $\check{\theta}$ that does not depend on Q and is a consistent estimator of θ^* . For example, the initial estimate $\check{\theta}$ could be the one optimizing the same function (3) by setting $\hat{Q}_t = 0$, i.e., without augmentation. Similar to (9), the weighted importance sampling is used for numerical stability. Then the initial estimator $\check{\theta}$ is found by

$$\check{\theta} = \arg \min_{\theta} \left\{ -\frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \frac{\rho_{i,1:t}^{\theta, \mu}}{\omega_{1:t}^{\theta, \mu}} R_{i,t} + \lambda_{\theta} \|\theta\|_1 \right\} \quad \text{s.t.} \quad \|\theta\|_2 \leq 1. \quad (10)$$

This estimate is consistent when the behavior policy μ is known and does not require estimation of the nuisance parameters.

Based on this initial estimate $\check{\theta}$, we model the nuisance parameter Q using a linear model of basis functions for each stage separately. Assume that $\phi(\mathbf{x}_t)$ is a basis function of the covariates \mathbf{x}_t at some stage t of dimension d' and it includes an intercept. In practice, $\phi(\mathbf{x}_t)$ can be taken to be the linear function, polynomial function, Gaussian radial basis functions, splines, wavelet basis, etc. of \mathbf{x}_t (Luckett et al., 2020; Shi et al., 2022). Let $\Phi(\mathbf{X}_t, A_t) := [\phi^T(\mathbf{X}_t), A_t \cdot \phi^T(\mathbf{X}_t)]^T$ and we can fit the model $Q_t(\mathbf{X}_t, A_t) = \Phi(\mathbf{X}_t, A_t)^T \beta_t$ at stage t , where β_t is in the dimension $2d'$. We discuss three different ways for defining the nuisance parameter Q .

In the first method, we define $Q_t^{\theta}(\mathbf{x}_t, a_t) := \mathbb{E}^{\theta}[\sum_{k=t}^T R_k | \mathbf{X}_t = \mathbf{x}_t, A_t = a_t]$ as the Q-function at stage t . Based on the tower property of conditional expectation, we have $Q_t^{\theta}(\mathbf{x}_t, a_t) := \mathbb{E}^{\theta}[R_t + Q_{t+1}^{\theta}(\mathbf{X}_{t+1}, A_{t+1}) | \mathbf{X}_t = \mathbf{x}_t, A_t = a_t]$. Following the idea of Q-learning (Zhu et al., 2019), we can estimate $\hat{\beta}_T, \dots, \hat{\beta}_1$ backwardly by plugging in the initial $\check{\theta}$. Denote the estimated Q-function at stage t for policy $\pi^{\check{\theta}}$ as

$$q_{i,t}^{\check{\theta}} = R_{i,t} + \sum_{a \in \mathcal{A}} \pi^{\check{\theta}}(a | \mathbf{X}_{i,t+1}) \Phi(\mathbf{X}_{i,t+1}, a)^T \hat{\beta}_{t+1}, \quad (11)$$

where $\hat{\beta}_{T+1}$ is taken as the zero vector. To find a sparse estimate of β_t , we minimize the square loss function with the L_1 penalty, i.e.,

$$\hat{\beta}_t^{(1)} = \arg \min_{\beta_{i,t} \in \mathbb{R}^{2d'}} \frac{1}{n} \sum_{i=1}^n \left\{ q_{i,t}^{\check{\theta}} - \Phi(\mathbf{X}_{i,t}, A_{i,t})^T \beta_t \right\}^2 + \lambda_{\beta_t^{(1)}} \|\beta_t\|_1, \quad (12)$$

where $\lambda_{\beta_t^{(1)}}$ is a tuning parameter.

The second method is a heuristic algorithm to minimize the variance of the estimator $\hat{V}_t(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})$ of the value function at stage t , where

$$\hat{V}_t(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}) = \frac{1}{n} \sum_{i=1}^n \left\{ \rho_{i,1:t}^{\boldsymbol{\theta}, \boldsymbol{\mu}} [R_{i,t} - \hat{Q}_t(\mathbf{X}_{i,t}, A_{i,t})] + \rho_{i,1:(t-1)}^{\boldsymbol{\theta}, \boldsymbol{\mu}} \hat{U}_t(\mathbf{X}_{i,t}) \right\}.$$

Note that the variance can be decomposed as $\text{Var}(\hat{V}_t(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})) = \mathbb{E}\{\hat{V}_t(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})\}^2 - \{\mathbb{E}\hat{V}_t(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})\}^2$. Since $\mathbb{E}\hat{V}_t(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}) = \mathbb{E}^{\boldsymbol{\theta}} R_t$ for any $\hat{\boldsymbol{\eta}}$, the second term is irrelevant to the nuisance parameter $\hat{\boldsymbol{\eta}}$. For computational efficiency, we only need to minimize $\mathbb{E}\{\hat{V}_t(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}})\}^2$ using the sample average. With initial estimate $\check{\boldsymbol{\theta}}$ and the weighted importance sampling, we have

$$\begin{aligned} \hat{\boldsymbol{\beta}}_t^{(2)} = \arg \min_{\boldsymbol{\beta}_{i,t} \in \mathbb{R}^{2d'}} \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\rho_{i,1:t}^{\check{\boldsymbol{\theta}}, \boldsymbol{\mu}}}{\omega_{1:t}^{\check{\boldsymbol{\theta}}, \boldsymbol{\mu}}} R_{i,t} - \right. \\ \left. \left[\frac{\rho_{i,1:t}^{\check{\boldsymbol{\theta}}, \boldsymbol{\mu}}}{\omega_{1:t}^{\check{\boldsymbol{\theta}}, \boldsymbol{\mu}}} \Phi(\mathbf{X}_{i,t}, A_{i,t}) - \frac{\rho_{i,1:(t-1)}^{\check{\boldsymbol{\theta}}, \boldsymbol{\mu}}}{\omega_{1:(t-1)}^{\check{\boldsymbol{\theta}}, \boldsymbol{\mu}}} \sum_{a \in \mathcal{A}} \pi^{\check{\boldsymbol{\theta}}}(a | \mathbf{X}_{i,t}) \Phi(\mathbf{X}_{i,t}, a) \right]^T \boldsymbol{\beta}_t \right\}^2 + \lambda_{\beta_t^{(2)}} \|\boldsymbol{\beta}_t\|_1, \end{aligned} \quad (13)$$

where $\lambda_{\beta_t^{(2)}}$ is a tuning parameter.

Note that both methods are essentially minimizing the asymptotic variance of the value function estimator. While the first method estimates Q_t as the Q-function of $\pi^{\boldsymbol{\theta}^*}$, which achieves semiparametric efficiency (Kallus and Uehara, 2020), the second method directly minimizes the finite sample variance, which will converge to the asymptotic variance. Furthermore, (12) and (13) both minimize the square loss with an L_1 penalty, and can be easily solved by existing Lasso packages with properly constructed predictors and responses. To improve the finite sample performance, we refit $\hat{\boldsymbol{\beta}}_t$ on its nonzero components for all t using linear regression. A refitted Lasso estimator means that we re-estimate the parameter on the support of the original Lasso estimator using the original loss function without the L_1 penalty. It has been shown that a refitted Lasso estimator usually leads to better finite sample performance than the original Lasso estimator (Zhang and Zhang, 2014; Ning and Liu, 2017). This refitted estimator may be less biased and less sensitive to the choice of tuning parameters of the penalty.

In the naive method, we can take $\lambda_{\beta_t^{(0)}} = \mathbf{0}$, so that $Q_t = 0$ for all t . Then the estimated nuisance parameter Q_t at stage t is

$$\hat{Q}_t^{(m)}(\mathbf{X}_t, A_t) = \Phi(\mathbf{X}_{1:t}, A_{1:t})^T \hat{\boldsymbol{\beta}}_t^{(m)} \quad (14)$$

for methods $m = 0, 1, 2$.

Optimizing constrained nonconvex nondifferentiable functions. To solve for $\check{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\theta}}$, note that (10) and (9) are both constrained nonconvex nondifferentiable optimization problems. To deal with the L_1 penalty, we use the proximal coordinate descent algorithm (Sra et al., 2011). To ensure that the L_2 norm of the estimated parameter is bounded by one, we normalize the parameter by its L_2 norm in each iteration of the coordinate descent. Since nonconvex problems may converge to a local minimizer, we start the optimization

Algorithm 2: Implementation for the optimization and inference of high-dimensional MSTP parameters

Input : n trajectories $\mathbf{D}_{1:n}$, the number of bootstrap iterations B , the confidence level $1 - \alpha$, the method m for estimating the nuisance parameter.

Output: The sparse estimator $\hat{\boldsymbol{\theta}}$, the one-step estimator $\tilde{\boldsymbol{\theta}}_j$ with its CI for $j \in \{1 : d\}$.

- 1 Obtain an initial estimator $\check{\boldsymbol{\theta}}$ by solving (10) using Algorithm 4, where $\lambda_{\check{\boldsymbol{\theta}}}$ is tuned by cross-validation;
- 2 Estimate the nuisance parameter $\hat{\mathbf{Q}}^{(m)}$ with $\check{\boldsymbol{\theta}}$ using Algorithm 3, where $\lambda_{\beta_1}, \dots, \lambda_{\beta_T}$ are tuned by cross-validation at each stage $t \in \{1 : T\}$;
- 3 Obtain the sparse estimator $\hat{\boldsymbol{\theta}}$ by solving (9) with $\hat{\mathbf{Q}}^{(m)}$ using Algorithm 4, where $\lambda_{\hat{\boldsymbol{\theta}}}$ is tuned by cross-validation;
- 4 **for** $j \in \{1 : d\}$ **do**
- 5 Calculate the one-step estimator $\tilde{\boldsymbol{\theta}}_j$ following the steps in (4), (5), and (6), where $\lambda_{\mathbf{w}_j}$ is tuned by cross-validation, and the gradient and Hessian matrix are estimated numerically by (17) and (21);
- 6 **end**
- 7 **for** $b \in \{1 : B\}$ **do**
- 8 Randomly draw a bootstrap sample of size n ;
- 9 Obtain the bootstrap estimator $\hat{\boldsymbol{\theta}}_b$ by refitting (9) with the bootstrap sample, $\hat{\mathbf{Q}}^{(m)}$ and $\lambda_{\hat{\boldsymbol{\theta}}}$ using Algorithm 4;
- 10 **for** $j \in \{1 : d\}$ **do**
- 11 Calculate the one-step estimator $\tilde{\boldsymbol{\theta}}_{b,j}$ following the steps in (4), (5), and (6) with $\lambda_{\mathbf{w}_j}$ and the numerically estimated gradient and Hessian matrix;
- 12 **end**
- 13 **end**
- 14 **for** $j \in \{1 : d\}$ **do**
- 15 Obtain the $(1 - \alpha) \times 100\%$ CI of $\tilde{\boldsymbol{\theta}}_j$ by finding the $\alpha/2$ and $1 - \alpha/2$ quantiles of $\{\tilde{\boldsymbol{\theta}}_{b,j}\}_{b=1}^B$.
- 16 **end**

from different starting points for better numerical results. Similar to the procedure for estimating the nuisance parameters, we refit $\check{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\theta}}$ on their support using the trust-region constrained algorithm, which is suitable for minimization problems with constraints. Although we require $\|\boldsymbol{\theta}\|_2 \leq 1$, the estimated $\hat{\boldsymbol{\theta}}$ actually has $\|\hat{\boldsymbol{\theta}}\|_2 = 1$ to approximate a deterministic policy. The optimization algorithm is summarized in Algorithm 4.

The full implementation algorithm for finding the estimate and the CI of high-dimensional policy parameters is summarized in Algorithm 2. Note that all the tuning parameters are selected using the full dataset and then fixed during bootstrap.

3. Theoretical Results

Denote $s_{\boldsymbol{\xi}} := \|\boldsymbol{\theta}^*\|_0$ and $s_{\mathbf{w}_j} := \|\mathbf{w}_j^*\|_0$ to be the numbers of nonzero elements in the corresponding vectors. We assume the following assumptions hold for the variables and the convergence rate of the nuisance parameters.

Assumption 4 *Assume the rewards R_t are bounded in the sense that $\|R_t\|_{\infty} \leq r$ for some $r > 0$ and for all $t \in \{1 : T\}$. Assume the covariates are bounded such that $\|\mathbf{X}_t\|_{\infty} \leq z$ and $|\mathbf{v}_j^{*T} \mathbf{X}_t| \leq z$ for some $z > 0$ and for all $t \in \{1 : T\}$. Suppose that there exists some constant $u > 0$ such that $|\theta_0^*| \geq u$.*

Assumption 5 *We assume that $(s_{\boldsymbol{\xi}} + s_{\mathbf{w}_j}) \log d / \sqrt{n} \rightarrow 0$. In addition, suppose that $\hat{\boldsymbol{\eta}} \in \mathcal{H}_n$ with probability no less than $1 - \Delta_n$, where*

$$\mathcal{H}_n := \{\boldsymbol{\eta} \in \mathcal{H} : \|\eta_t\|_{\mathbb{P}, \infty} \leq r, \|\eta_t - \bar{\eta}_t\|_{\mathbb{P}, 2} \leq \delta_n \text{ for all } t\},$$

and $\delta_n = o(1)$, $\Delta_n = o(1)$ are positive constants.

Assumption 6 *Assume $V(\boldsymbol{\theta})$ is κ -strongly concave at $\boldsymbol{\theta}^*$. That is, there exists $\kappa > 0$ such that $\langle \nabla_{\boldsymbol{\xi}} V(\boldsymbol{\theta}^*), \Delta_{\boldsymbol{\xi}} \rangle - [V(\boldsymbol{\theta}^* + \Delta_{\boldsymbol{\theta}}) - V(\boldsymbol{\theta}^*)] \geq \frac{\kappa}{2} \|\Delta_{\boldsymbol{\xi}}\|_2^2$ for all $\Delta_{\boldsymbol{\xi}} \in \mathbb{B}(R)$ for some radius R , where $\mathbb{B}(R)$ is the ball with radius R defined by the L_2 norm, and $\Delta_{\boldsymbol{\theta}} = (\pm \sqrt{1 - \|\Delta_{\boldsymbol{\xi}}\|_2^2}, \Delta_{\boldsymbol{\xi}})$ is the concatenated error bound for $\boldsymbol{\theta}$.*

Assumption 4 requires the boundedness of the variables. The conditions about the covariates \mathbf{X}_t follow the example of generalized linear models in Ning and Liu (2017). We assume that the random policy that assigns both actions with equal probabilities will not be the optimal policy. Without loss of generality, we assume the intercept is nonzero to achieve this. If the intercept is zero, we can shift the j th feature by a constant u/θ_j for some coordinate j where $\theta_j \neq 0$.

Assumption 5 deals with the convergence rate of the nuisance parameters. Since we assume that we are using data from an MRT and $\boldsymbol{\mu}$ is known, we do not need the model for Q_t to be correctly specified. Moreover, \hat{Q}_t can converge to its limit at any rate. This assumption can be easily satisfied by almost all learning methods, including regularized methods like Lasso, ridge regression, or elastic net.

Assumption 6 is used to verify the restricted strong convexity, which is one of the sufficient conditions for proving the convergence of parameters regularized by the L_1 penalty. The local concavity at $\boldsymbol{\theta}^*$ is empirically verified with numerical examples in Appendix B.4.

Lemma 1 *Under Assumptions 1-6, when $\lambda_{\boldsymbol{\theta}} \simeq \sqrt{\log d/n}$, we have*

$$\|\hat{\boldsymbol{\xi}} - \boldsymbol{\xi}^*\|_1 = \mathcal{O}_{\mathbb{P}}(s_{\boldsymbol{\xi}} \sqrt{\log d/n}). \quad (15)$$

Lemma 1 shows that the sparse estimator $\hat{\boldsymbol{\xi}}$, which is a permutation of $\hat{\boldsymbol{\theta}}$, is a consistent estimator of $\boldsymbol{\xi}^*$. We will then show that the one-step estimator $\hat{\theta}_j$ is a consistent estimator of θ_j^* with a \sqrt{n} convergence rate.

Assumption 7 *Suppose that the covariance matrix $\boldsymbol{\Sigma}^* = \text{Var}[\nabla_{\boldsymbol{\xi}} l_0(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]$ is positive definite and finite.*

Assumption 7 is used in the multivariate central limit theorem to prove the asymptotic normality of the score function. We follow the proof of Ning and Liu (2017) to decorrelate ν_j from the parameter θ_j that we are interested in, and the assumptions will be used to verify the conditions in Theorem 3.2 (Ning and Liu, 2017). The main challenge lies in the nuisance parameters η , which need to be estimated and are also high-dimensional. We will use the technique in Chernozhukov et al. (2018a) to decorrelate the nuisance parameters.

Theorem 2 *Under Assumptions 1-7, the one-step estimator $\tilde{\theta}_j$ satisfies*

$$\sqrt{n}(\tilde{\theta}_j - \theta_j^*) I_{\theta_j|\nu_j}^* / \sigma_j^{*1/2} \Rightarrow N(0, 1) \quad (16)$$

for $j \in \{1 : d\}$, where σ_j^* is defined in (7).

Theorem 2 shows that the plug-in one-step estimator is asymptotically normal. Since $\hat{I}_{\theta_j|\nu_j}$ is consistent for $I_{\theta_j|\nu_j}^*$ and $\hat{\sigma}_j$ is consistent for σ_j^* , we can construct CIs as in (8) based on the theorem. Note that σ_j^* depends on the nuisance parameter η . Therefore, although the limit and the convergence rate of $\hat{\eta}$ do not affect the convergence rate of $\tilde{\theta}_j$, the limit η does influence the asymptotic variance of $\tilde{\theta}_j$.

4. Simulation Study

In this section, we test our proposed algorithm for the optimization and inference of low-dimensional parameters in high-dimensional settings in two simulated scenarios.

Assume that the data are from an MRT, and the action $A_{i,t}$ takes values from $\{-1, 1\}$ with equal probability at each stage t for all individuals i . The variable $X_{i,t,j}$ represents the j th variable of individual i at stage t . The first two variables for $j = 1, 2$ are important variables that influence the reward, while the other variables for $j \geq 3$ are noise variables. The initial states are generated from a multivariate normal distribution, i.e., $\mathbf{X}_{i,1} \sim N(\mathbf{0}, \Sigma)$ for each individual $i \in \{1 : n\}$. In the covariance matrix Σ , the variance of variable j is $\Sigma_{j,j} = 4$ for all $j \in \{1 : d\}$. To allow dependence between high-dimensional features, the covariance between the important variables is $\Sigma_{1,2} = \Sigma_{2,1} = 1$, and the covariance between variables $j - 1$ and j is $\Sigma_{j-1,j} = \Sigma_{j,j-1} = 0.2$ for $j \in \{3 : d\}$. Other entries in Σ take a value of zero. For stages $t \in \{2 : (T + 1)\}$, the variables and the rewards are generated as follows. In Scenario 1,

$$\begin{aligned} X_{i,t,1} &= 0.8\tilde{X}_{i,t-1,1} + 0.3A_{i,t-1}\tilde{X}_{i,t-1,1} + 0.2\tilde{X}_{i,t-1,2} + \epsilon_{i,t,1}, \\ X_{i,t,2} &= 0.8\tilde{X}_{i,t-1,2} - 0.3A_{i,t-1}\tilde{X}_{i,t-1,2} + A_{i,t-1} \tanh\{(\tilde{X}_{i,t-1,1} - \tilde{X}_{i,t-1,2})/2\} + \epsilon_{i,t,2}, \\ X_{i,t,j} &= 0.9\tilde{X}_{i,t-1,j} + \epsilon_{i,t,j}, \quad j \in \{3 : d\}, \\ R_{i,t-1} &= \log\{1 + \exp(X_{i,t,1} + X_{i,t,2})\} - 0.5A_{i,t}, \end{aligned}$$

and in Scenario 2,

$$\begin{aligned} X_{i,t,1} &= 0.8\tilde{X}_{i,t-1,1} + 0.3A_{i,t-1}\tilde{X}_{i,t-1,1} + 0.1\tilde{X}_{i,t-1,2} + \epsilon_{i,t,1}, \\ X_{i,t,2} &= 0.8\tilde{X}_{i,t-1,1} + 0.2A_{i,t-1}\tilde{X}_{i,t-1,2} + 0.2\tilde{X}_{i,t-1,2} + \epsilon_{i,t,2}, \\ X_{i,t,j} &= 0.9\tilde{X}_{i,t-1,j} + \epsilon_{i,t,j}, \quad j \in \{3 : d\}, \\ R_{i,t-1} &= X_{i,t,1} + X_{i,t,2} - 0.6A_{i,t}, \end{aligned}$$

for $i \in \{1 : n\}$. Here, $\tilde{X}_{i,t,j}$ is the sequence of exponentially weighted moving averages of $X_{i,t,j}$ such that $\tilde{X}_{i,1,j} = X_{i,1,j}$ and $\tilde{X}_{i,t,j} = 0.2\tilde{X}_{i,t-1,j} + 0.8X_{i,t,j}$ for all $j \in \{1 : d\}$ and $t \geq 2$. Let $\epsilon_{i,t,j} \stackrel{i.i.d.}{\sim} N(0, 0.4^2)$ for all i, j , and $t \geq 2$. Under these scenarios, the Markov assumption is violated since the reward $R_{i,t}$ is a function of $\mathbf{X}_{i,t+1}$, which depends on the states in all previous stages.

In this simulation, we experiment with three different constructions of the nuisance parameter \mathbf{Q} as discussed in Section 2.3. The number of stages T is taken to be 1, 5, or 10. The dimension is fixed at $d = 50$. The sample size n ranges among 200, 400, 800, 1600, and 3200. In our simulation, the scaling parameter τ is fixed at 0.2, since our experiments show that a larger τ leads to a stochastic policy far from the true optimal deterministic policy, and a smaller τ may cause unstable computation. We take the number of bootstraps to be $B = 100$. The value function of each estimated policy is calculated based on an independent test set of size 200,000 generated by this policy. We repeat the whole procedure $W = 100$ times for each scenario.

The simulation results for Scenario 1 and Scenario 2 are presented in Figures 1 and 2, respectively. We report the average reward for the policy $\pi^{\hat{\theta}}$ based on the sparse estimator $\hat{\theta}$, and the mean absolute deviation (MAD) and the coverage probability of the one-step estimator $\tilde{\theta}$. We find the true minimizer θ^* of the loss function within the class Π by grid-search for each T . In particular, we estimate the value function on an independent test set of size 200,000 for θ on the grids inside the unit ball. The true value of θ^* in different scenarios is reported in Appendix B.4.

We compare our proposed method with Penalized Efficient Augmentation and Relaxation Learning (PEARL) (Liang et al., 2022) when $T = 1$. PEARL is a method for estimating the optimal Individualized Treatment Rule (ITR) and conducting statistical inference from high-dimensional data in single-stage decision problems. It utilizes the data-splitting method to account for the slow convergence rate of nuisance parameter estimations. In addition, it follows the inference procedure in Ning and Liu (2017), first finding a sparse estimator and then obtaining the one-step estimator. We use the package `ITRInference` for implementation (Liang et al., 2022). Since an ITR generated by PEARL takes the selected treatment with probability one, any positive multiplication of $\hat{\theta}_{\text{PEARL}}$ will yield the same ITR. However, due to the strict convexity of their loss function, the estimated optimal parameter $\hat{\theta}_{\text{PEARL}}$ is always unique. Therefore, the package does not impose requirements on the scale of the parameters. Their simulation study estimated the coverage probability of the constructed CI based on the limit of parameter estimates, which may not be identical to the true parameters for data generation. In practice, we rescale all parameters so that the average one-step estimator $1/W \sum_{w=1}^W \tilde{\theta}_{\text{PEARL},w,1}$ at coordinate $j = 1$ matches θ_1^* at $j = 1$, where θ^* is the true optimal parameter of a deterministic policy.

For a better comparison across different values of T , we report the average reward $\bar{V}(\theta) = \frac{1}{T}V(\theta)$ instead of the value function $V(\theta)$. The first rows in Figures 1 and 2 present the mean of $\bar{V}(\hat{\theta})$ and its 95% CI based on the standard error across W replications. The sparse estimator $\hat{\theta}$ constructed using different nuisance parameters $Q^{(0)}, Q^{(1)}, Q^{(2)}$ is compared with the random policy used to generate the data. We observe that in both scenarios, the nuisance parameter $Q^{(1)}$ yields the largest average reward and the smallest standard error, compared to the other two nuisance parameters and the random policy. Although

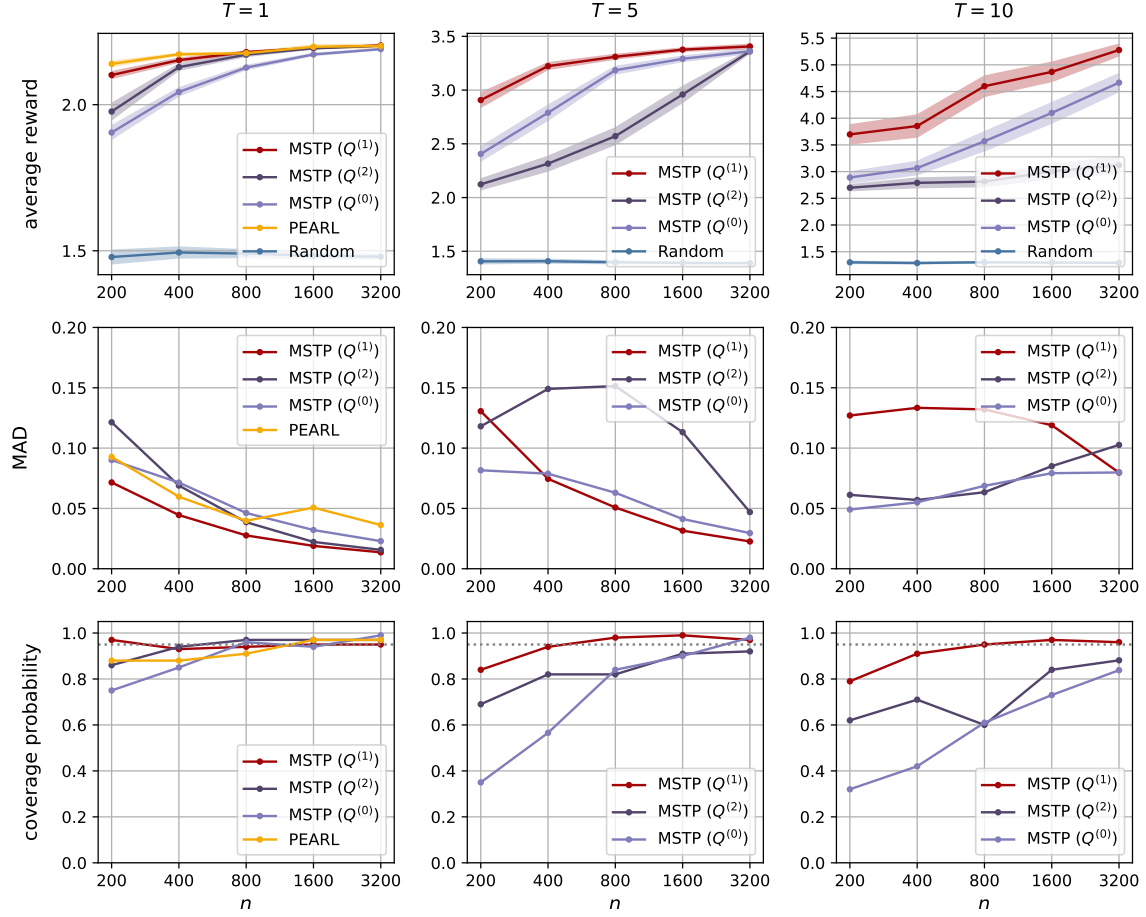


Figure 1: The first row presents the average reward of MSTP, PEARL, and the random policy in Scenario 1. The second and third rows show the MAD and coverage probability of θ_1 . The columns correspond to different values of T . The dotted line indicates the nominal coverage probability for θ_1 .

both $Q^{(1)}$ and $Q^{(2)}$ essentially minimize the asymptotic variance of the value function estimator, $Q^{(2)}$ depends on the importance sampling weight and becomes numerically unstable as T increases. When $T = 1$, we further compare with PEARL, since it is only applicable in single-stage decision problems. PEARL estimates a deterministic policy, which yields a higher average reward in Scenario 1. Although the proposed method estimates a stochastic policy, it is stable when $T = 1$ and yields a larger average reward in Scenario 2.

We further report the MAD and the coverage probability of the important variable θ_1 in the second and third rows of Figures 1 and 2. The MAD and the coverage probability of the second important variable θ_2 and the noise variables $\theta_{3:d}$ are reported in Appendix B.1. For the one-step estimator of PEARL, the MAD is approximated as $\tilde{\sigma}_{\text{PEARL}}/1.4826$, where $\tilde{\sigma}_{\text{PEARL}}$ is the estimated standard deviation in the asymptotically normal distribution of the one-step estimator. We observe that the coverage probabilities of $\tilde{\theta}$ are concentrated near the nominal coverage of 95% when the sample size is large enough. When T is small,

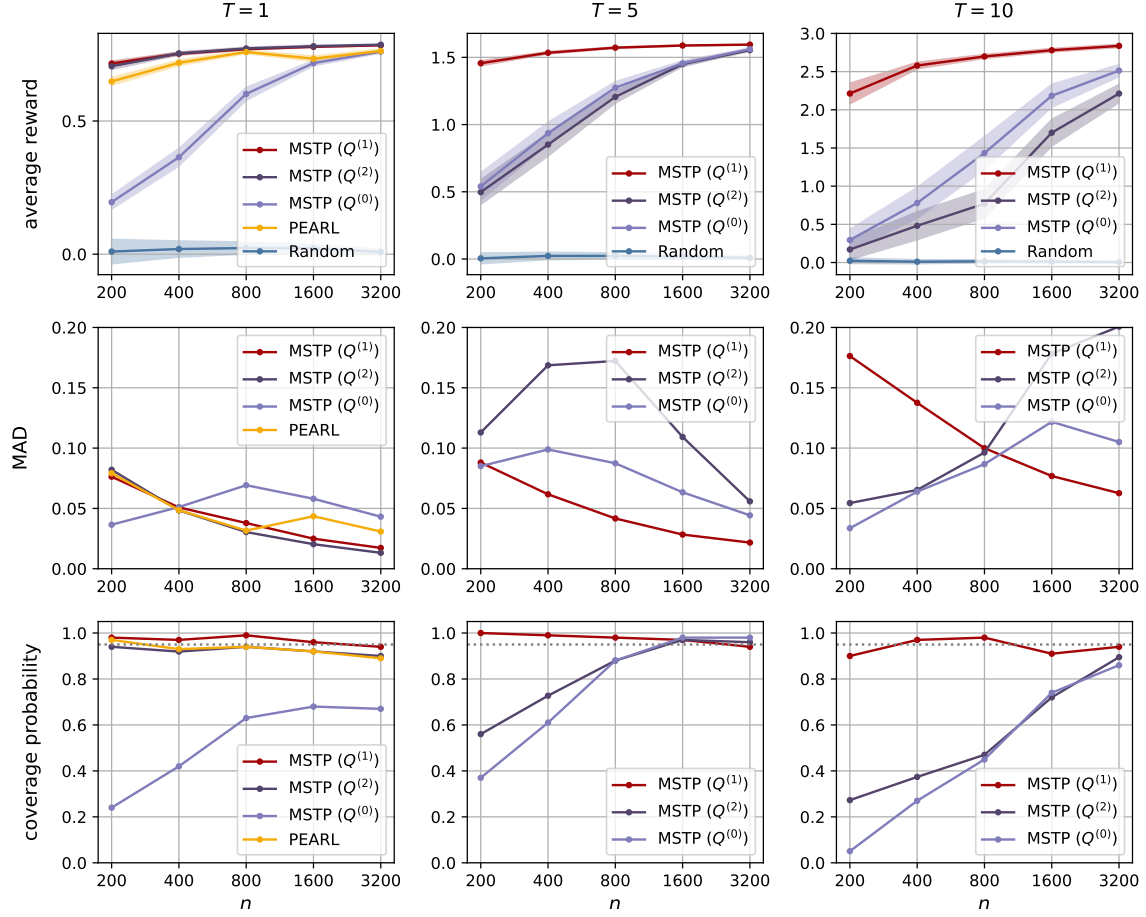


Figure 2: The first row presents the average reward of MSTP, PEARL, and the random policy in Scenario 2. The second and third rows show the MAD and coverage probability of θ_1 . The columns correspond to different values of T . The dotted line indicates the nominal coverage probability for θ_1 .

the CIs achieve the nominal coverage probability even with a small sample size. However, the difficulty of policy learning and inference increases as T increases. More samples are needed to learn the optimal policy and reach the nominal coverage probability. The MAD generally decreases with the increasing sample size if the sparse estimator is well estimated. The MAD of MSTP with $Q^{(2)}$ or $Q^{(0)}$ may be small when $n = 200, 400$ for $j = 1, 2$ since $\hat{\theta}_1$ or $\hat{\theta}_2$ are estimated to be zero on a large portion of bootstrap samples, but the one-step estimator cannot correct the bias with a small sample size. The MAD of $\tilde{\theta}$ based on $Q^{(1)}$ is smaller than that of the other two nuisance parameters and PEARL, while achieving a coverage probability even closer to the nominal value. Therefore, the above experiments demonstrate the influence of the nuisance parameters on the asymptotic efficiency, and the advantage of using $Q^{(1)}$ in the doubly robust estimator.

5. Real Data Analysis

The OhioT1DM dataset is a collection of data from individuals with T1D collected through continuous glucose monitors (CGMs), insulin usage, and self-reported events (Marling and Bunesco, 2020). It is designed to support research in blood glucose prediction, diabetes management, and related healthcare applications. In the 2018 cohort, 6 patients participated in an 8-week study. We apply the proposed method to the OhioT1DM dataset to learn a treatment policy for patients with T1D.

After excluding extremely rare events, we use 12 features to construct the MSTP, including the current glucose level, the long-acting basal insulin level, carbohydrate estimate for a meal, and others. Since the glucose level, heart rate, galvanic skin response (GSR), skin temperature, air temperature, and step counts are recorded every 5 minutes using CGM, the treatment decision is defined as whether to deliver short-acting bolus insulin every 5 minutes. The reward is defined as a function of the glucose level in the next 5 minutes, following Shi et al. (2020b). We take $T = 12$, so that the value function is the sum of rewards over one hour. The detailed description of the features and the data cleaning methods are provided in Appendix C.

Based on the results in Section 4, we use the nuisance parameter $Q^{(1)}$ in Algorithm 2. In our analysis, the propensity score is estimated using logistic regression with an L_1 penalty. All features have been standardized. The reward is standardized before the analysis but is returned to the original scale in the following results.

To estimate the mean and standard error of the average reward, we merge the training and test examples in the original dataset and obtain a combined dataset of 7824 trajectories. We then randomly split the combined dataset into training and test sets with probabilities 0.8 and 0.2 for $W = 100$ replications. The average reward of the behavior policy is estimated by taking the average of all rewards in the training set. The average reward of the estimated policy is calculated using the IPW estimator with normalized weights on the test set. After applying Algorithm 2 for policy learning, the average reward for the behavior policy is -6.260 , with a standard error of 0.008 ; the average reward for the estimated policy is -2.367 , with a standard error of 0.100 . This suggests that the estimated policy could significantly improve upon the behavior policy.

We use the training and test set splitting in the original dataset to construct the policy using the sparse estimator $\hat{\theta}$ and to construct the CIs of $\tilde{\theta}$, which are reported in Table 1. Based on the results in Table 1, the current glucose level, meal, hypoglycemic episode, heart rate, and GSR are selected, but only the glucose level is significant. The estimated policy suggests delivering bolus insulin when the current glucose level is high and avoiding it when the glucose level is already low. Since “hypoglycemic episode” indicates low glucose levels, one might expect its coefficient to be negative, but the estimated coefficient is positive. This discrepancy arises because hypoglycemic episodes are self-reported, reflecting patients’ awareness of the condition and prompting them to consume meals or snacks immediately. Out of 78 self-reported events in the entire dataset, patients consumed a meal in 74 events. As a result, “hypoglycemic episode” becomes an indicator of an increase in blood glucose levels in subsequent times. Nevertheless, the coefficient is not significant based on its CI, as this variable is highly correlated with glucose levels.

j	Feature	$\hat{\theta}_j$	$\tilde{\theta}_j$	CI of $\tilde{\theta}_j$
0	Intercept	0.000		
1	Basal	0.000	-0.001	(-0.354, 0.039)
2	Glucose level	0.872	0.870	(0.302, 0.991)
3	Meal	0.463	0.460	(-0.032, 0.553)
4	Sleep	0.000	-0.001	(-0.093, 0.045)
5	Work	0.000	-0.001	(-0.057, 0.060)
6	Hypoglycemic episode	0.084	0.081	(-0.041, 0.113)
7	Exercise	0.000	-0.000	(-0.061, 0.290)
8	Heart rate	0.102	0.097	(-0.196, 0.171)
9	GSR	0.084	0.079	(-0.194, 0.547)
10	Skin temperature	0.000	-0.002	(-0.119, 0.085)
11	Air Temperature	0.000	-0.002	(-0.095, 0.091)
12	Steps	0.000	-0.000	(-0.150, 0.772)

Table 1: The estimated sparse policy parameter $\hat{\theta}_j$ and the CIs of the one-step estimator $\tilde{\theta}_j$.

6. Discussion

In this work, we focus on the multi-stage decision problem and propose a method for estimating the high-dimensional MSTP and the CIs of its parameters. We first estimate the MSTP based on the AIPWE of the value function with an L_1 penalty to encourage sparsity and an L_2 constraint to avoid unboundedness. Then, we compute the one-step estimators, which are asymptotically normal and suitable for statistical inference. The proposed one-step estimator is shown to achieve nominal coverage probabilities in simulation studies. While the choice of nuisance parameter estimation does not affect the convergence rate of the low-dimensional policy parameter, it influences the asymptotic efficiency. We compare different estimators \hat{Q} in the simulation study and show that Q_t constructed as the Q-function generates the highest value function and the smallest MAD of the estimated policy parameters. Our code for the simulation study and real data analysis can be accessed at <https://github.com/DqGao/MSTP>.

We assume that the behavior policy μ is known, as in MRTs. When data from an observational study are used for learning, μ is unknown and needs to be estimated. Our framework naturally extends to observational studies. The one-step estimator can still achieve the \sqrt{n} convergence rate as long as either the Q-function or the propensity score is correctly specified. Since we are most interested in the significance of the selected variables, we focus on constructing confidence intervals for each parameter separately. Similar methods can be used to construct confidence regions for multi-dimensional parameters (Ning and Liu, 2017), if the joint effect of multiple features is of interest. In this work, we focus on binary treatments. Extending the method to multiple treatments is an interesting open problem (Qi et al., 2020; Ma et al., 2022, 2023). A possible approach is to replace the sigmoid probability in the policy class Π with a softmax probability, but the conditions for valid statistical inference require further investigation.

For decision problems with a long horizon T , the ratio $\rho_{1:t}^{\theta, \mu}$ can become extremely unstable, especially when the feature dimension d is large. To handle this instability, an alternative method based on the marginalized distribution of the current state and action has been proposed in Liu et al. (2018a); Kallus and Uehara (2020); Liao et al. (2022). However, it requires the Markov assumption, which is often violated in mHealth applications (Shi et al., 2020b). Moreover, this approach requires either the estimation of a marginal distribution for each stage t , or the convergence of the offline data to the stationary distribution (Kallus and Uehara, 2020; Liao et al., 2022). Particularly, estimating the marginal distribution becomes challenging with high-dimensional continuous features. Due to the curse of dimensionality (Hastie et al., 2009), the data for density ratio estimation can be very sparse. Lockett et al. (2020) only uses the importance sampling weight at one decision stage based on the Bellman optimality equation, but it requires the value function to be estimated at each step of the gradient descent algorithm with respect to the policy parameters. Another approach under the Markov assumption is to use partial history importance weighting with assumptions on mixing and overlap between the behavior and target policies (Hu and Wager, 2023). If the Markov chain under the target policy mixes quickly, the policy before mixing has little impact on later rewards. Thus, only the importance sampling weight after mixing is needed, efficiently reducing the variance of the value function estimator for a long horizon.

Acknowledgment

This research was supported by NIH grants R01 GM124104 and R01 HL125440, and NSF grants DMS 2100729, SES 2217440, and DMS 2515561.

Appendix A. Implementation Details

In this section, we present the implementation details omitted in Section 2.3.

A.1 Estimating the Nuisance Parameters

In Section 2.3, we discussed three different methods for estimating the nuisance parameter Q . In Algorithm 3, we present the full algorithm for estimating $\hat{Q}^{(m)}$ for $m = 0, 1, 2$. The tuning parameter $\lambda_{\beta_t}^{(m)}$ is selected as the one that minimizes the mean squared error in (12) or (13) on the validation set. While for the method $m = 1$, the parameters $\beta_t^{(1)}$ must be estimated backwardly from the last stage to the first one, the parameters $\beta_t^{(2)}$ for the method $m = 2$ do not require a specific estimation order.

A.2 Optimization with L_1 Penalty and L_2 Constraint

We present the optimization method for a nonconvex loss function with an L_1 penalty and L_2 constraint as described in Section 2.3 in Algorithm 4. Notice that to find the global minimizer of a nonconvex function, we need to try different starting points. Since running the full algorithm from multiple starting points is computationally heavy, we compare the function values of multiple points for each coordinate separately at the beginning.

Algorithm 3: Estimate the Nuisance Parameter Q

Input : The method m , the initial estimate $\check{\theta}$, the tuning parameter $\lambda_{\beta_t}^{(m)}$.

Output: The estimated nuisance parameters $\hat{Q}^{(m)} = [\hat{Q}_{1:T}^{(m)}]$.

```

1 if  $m = 0$  then
2   | Take  $\lambda_{\beta_t^{(0)}} = 0$ ;
3 else if  $m = 1$  then
4   | for  $t = T, \dots, 1$  do
5   |   | Obtain  $q_{i,t}^{\check{\theta}}$  with  $\hat{\beta}_{t+1}$  by (11);
6   |   | Estimate  $\hat{\beta}_t^{(1)}$  by solving (12) with  $\check{\theta}$  and  $\lambda_{\beta_t}^{(1)}$  using a Lasso solver;
7   |   | Refit  $\hat{\beta}_t^{(1)}$  on its nonzero components using linear regression;
8   | end
9 else if  $m = 2$  then
10  | for each  $t \in \{1 : T\}$  do
11  |   | Estimate  $\hat{\beta}_t^{(2)}$  by solving (13) with  $\check{\theta}$  and  $\lambda_{\beta_t}^{(2)}$  using a Lasso solver;
12  |   | Refit  $\hat{\beta}_t^{(2)}$  on its nonzero components using linear regression;
13  | end
14 end
15 Find  $\hat{Q}_t^{(m)}$  by plugging  $\hat{\beta}_t^{(m)}$  into (14) for each  $t$ .
```

The tuning parameter $\lambda_{\check{\theta}}$ for the initial estimator $\check{\theta}$ or λ_{θ} for the sparse estimator $\hat{\theta}$ is selected as the one that maximizes the estimated value function $\hat{V}(\theta)$ on the validation set. Note that for $\lambda_{\check{\theta}}$, $Q_t = 0$ in $\hat{V}(\theta)$. We use the function `scipy.optimize.minimize` with `method='trust-constr'` in Python to solve the constrained minimization problem when refitting $\check{\theta}$ and $\hat{\theta}$ on their support.

A.3 Obtaining the One-Step Estimator

The tuning parameter λ_{w_j} is chosen for each feature j separately. It is selected as the value that minimizes the projection error in (4) on the validation set. To estimate the Dantzig-type estimator in (4), we use the package `cvxpy` to solve the constrained convex minimization problem.

A.4 Numerical Computation of the Gradient and Hessian matrix of the Loss Function

The direct estimation of the gradient and Hessian matrix of $\rho_{i,1:t}^{\theta,\mu}$ is unstable due to the probability product, which may cause even larger variability when constructing CIs for the one-step estimator and lead to undercoverage of the true parameter. Therefore, we propose to calculate them numerically. An estimate of the partial derivative at the coordinate $j \in \{1 : d\}$ is calculated using the symmetric difference quotient as

$$\hat{\nabla}_{\theta_j} \ell(\theta, \eta) = \frac{\ell(\theta^+, \eta) - \ell(\theta^-, \eta)}{2h_j}, \quad (17)$$

Algorithm 4: Optimization of a Loss Function with L_1 Penalty and L_2 Constraint

Input : A loss function ℓ to be minimized with L_1 penalty and L_2 constraint, the tuning parameter λ for the L_1 penalty, a starting point $\boldsymbol{\theta}^{(0)}$, the error threshold e , the maximum number of iterations M .

Output: The estimated parameter $\tilde{\boldsymbol{\theta}}$ or $\hat{\boldsymbol{\theta}}$.

```

1 for each coordinate  $j \in \{0 : d\}$  do
2   | Find the best starting point of the coordinate  $j$  using grid search;
3 end
4 Normalize  $\boldsymbol{\theta}^{(1)}$  by its  $\ell_2$ -norm;
5 Initialize the number of iterations  $m = 1$ ;
6 while  $m \leq M$  and  $\|\boldsymbol{\theta}^{(m)} - \boldsymbol{\theta}^{(m-1)}\|_2 \geq e$  do
7   |  $m \leftarrow m + 1$ ;
8   | for each coordinate  $j \in \{0 : d\}$  do
9     | Find the minimizer of  $\ell$  with respect to the current coordinate using the
      | BFGS algorithm
      |
      | 
$$\tilde{\theta}_j^{(m)} = \arg \min_{\theta_j} \ell((\theta_{0:(j-1)}^{(m)}, \theta_j, \theta_{(j+1):d}^{(m-1)}));$$

      |
      | Shrink  $\tilde{\theta}_j^{(m)}$  using soft-thresholding
      |
      | 
$$\theta_j^{(m)} = \text{sign}\{\tilde{\theta}_j^{(m)}\} \max(|\tilde{\theta}_j^{(m)}| - \lambda, 0);$$

      |
      | Normalize  $(\theta_{0:j}^{(m)}, \theta_{(j+1):d}^{(m-1)})$  by its  $\ell_2$ -norm;
10  | end
11 end
12 Re-estimate  $\tilde{\boldsymbol{\theta}}$  or  $\hat{\boldsymbol{\theta}}$  by minimizing the loss function  $\ell$  with the  $L_2$  constraint on the
    support of  $\boldsymbol{\theta}^{(m)}$  using the trust-region constrained algorithm;
```

where

$$h_j = \min \left\{ \frac{1}{\sqrt{nT}}, \sqrt{1 - \sum_{l \neq 0, j} \theta_l^2} - |\theta_j| \right\},$$

$$\theta_j^+ = \theta_j + h_j,$$

$$\theta_0^+ = \text{sign}\{\theta_0\} \sqrt{1 - \sum_{l \neq 0, j} \theta_l^2 - (\theta_j^+)^2}, \quad (18)$$

$$\theta_l^+ = \theta_l \quad \text{for } l \neq 0, j, \quad (19)$$

and $\boldsymbol{\theta}^-$ can be defined similarly with $\theta_j^- = \theta_j - h_j$. Here, $1/\sqrt{nT}$ is the common value for h_j used in numerical computations. However, since the ℓ_2 -norm is restricted to be 1 in our case, h_j sometimes needs to be even smaller depending on the initial estimate θ_j . If $\theta_0 = 0$, then $\sqrt{1 - \sum_{l \neq 0, j} \theta_l^2} - |\theta_j| = 0$. In this case, when $\theta_j \neq 0$, we estimate the gradient using

Newton's difference quotient as

$$\hat{\nabla}_{\theta_j} \ell(\boldsymbol{\theta}, \boldsymbol{\eta}) = \frac{\ell(\boldsymbol{\theta}^+, \boldsymbol{\eta}) - \ell(\boldsymbol{\theta}, \boldsymbol{\eta})}{-h_j \text{sign}\{\theta_j\}},$$

where

$$h_j = \min \left\{ \frac{1}{\sqrt{nT}}, 2|\theta_j| \right\},$$

$$\theta_j^+ = \theta_j - h_j \text{sign}\{\theta_j\},$$

and the other coordinates θ_l^+ for $l \neq j$ are calculated as in (18) and (19). When $\theta_0 = \theta_j = 0$, we use the gradient of $\boldsymbol{\theta}'$ as an approximation, where

$$\theta'_0 = 0, \quad \theta'_j = \frac{1}{\sqrt{nT}}, \quad \theta'_l = \theta_l \sqrt{1 - (\theta'_j)^2} \text{ for } l \neq 0, j. \quad (20)$$

Similarly, we use the symmetric difference quotient to find an estimate of the Hessian matrix:

$$\hat{\nabla}_{\theta_j \theta_k} \ell(\boldsymbol{\theta}, \boldsymbol{\eta}) = \frac{\ell(\boldsymbol{\theta}^{++}, \boldsymbol{\eta}) - \ell(\boldsymbol{\theta}^{+-}, \boldsymbol{\eta}) - \ell(\boldsymbol{\theta}^{-+}, \boldsymbol{\eta}) + \ell(\boldsymbol{\theta}^{--}, \boldsymbol{\eta})}{(2h_{jk})^2}. \quad (21)$$

When $j \neq k$,

$$h_{jk} = \min \left\{ \frac{1}{\sqrt{nT}}, \frac{1 - \sqrt{\sum_{l \neq 0} \theta_l^2}}{\sqrt{2}} \right\},$$

$$\theta_j^{++} = \theta_j + h_{jk}, \quad \theta_k^{++} = \theta_k + h_{jk},$$

$$\theta_0^{++} = \text{sign}\{\theta_0\} \sqrt{1 - \sum_{l \neq 0, j, k} \theta_l^2 - (\theta_j^{++})^2 - (\theta_k^{++})^2}, \quad (22)$$

$$\theta_l^{++} = \theta_l \text{ for } l \neq 0, j, k, \quad (23)$$

and $\boldsymbol{\theta}^{+-}, \boldsymbol{\theta}^{-+}, \boldsymbol{\theta}^{--}$ can be defined similarly with

$$\begin{aligned} \theta_j^{+-} &= \theta_j + h_{jk}, & \theta_k^{+-} &= \theta_k - h_{jk}, \\ \theta_j^{-+} &= \theta_j - h_{jk}, & \theta_k^{-+} &= \theta_k + h_{jk}, \\ \theta_j^{--} &= \theta_j - h_{jk}, & \theta_k^{--} &= \theta_k - h_{jk}. \end{aligned}$$

When $j = k$,

$$h_{jj} = \min \left\{ \frac{1}{\sqrt{nT}}, \frac{1}{2} \left(\sqrt{1 - \sum_{l \neq 0, j} \theta_l^2} - |\theta_j| \right) \right\},$$

$$\theta_j^{++} = \theta_j + 2h_{jj},$$

$$\theta_0^{++} = \sqrt{1 - \sum_{l \neq 0, j} \theta_l^2 - (\theta_j^{++})^2}, \quad (24)$$

$$\theta_l^{++} = \theta_l \text{ for } l \neq 0, j, k, \quad (25)$$

and $\boldsymbol{\theta}^{+-}, \boldsymbol{\theta}^{-+}, \boldsymbol{\theta}^{--}$ can be defined similarly with

$$\theta_j^{+-} = \theta_j^{-+} = \theta_j, \quad \theta_j^{--} = \theta_j - 2h_{jj}.$$

If $\theta_0 = 0$ and $\theta_j, \theta_k \neq 0$, we instead use Newton's difference quotient to estimate the Hessian matrix as

$$\hat{\nabla}_{\theta_j \theta_k} \ell(\boldsymbol{\theta}, \boldsymbol{\eta}) = \frac{\ell(\boldsymbol{\theta}^{++}, \boldsymbol{\eta}) - \ell(\boldsymbol{\theta}^{+-}, \boldsymbol{\eta}) - \ell(\boldsymbol{\theta}^{-+}, \boldsymbol{\eta}) + \ell(\boldsymbol{\theta}, \boldsymbol{\eta})}{(h_{jk} \text{sign}\{\theta_j\})(h_{jk} \text{sign}\{\theta_k\})}.$$

When $j \neq k$,

$$h_{jk} = \min \left\{ \frac{1}{\sqrt{nT}}, 2|\theta_j|, 2|\theta_k| \right\},$$

$$\theta_j^{++} = \theta_j - h_{jk} \text{sign}\{\theta_j\}, \quad \theta_k^{++} = \theta_k - h_{jk} \text{sign}\{\theta_k\},$$

and the other coordinates θ_l^{++} for $l \neq 0$ are calculated as in (22) and (23). Similarly, $\boldsymbol{\theta}^{+-}, \boldsymbol{\theta}^{-+}$ can be defined with

$$\theta_j^{+-} = \theta_j + h_{jk}, \quad \theta_k^{+-} = \theta_k,$$

$$\theta_j^{-+} = \theta_j, \quad \theta_k^{-+} = \theta_k + h_{jk}.$$

When $j = k$,

$$h_{jj} = \min \left\{ \frac{1}{\sqrt{nT}}, |\theta_j| \right\},$$

$$\theta_j^{++} = \theta_j - 2h_{jj} \text{sign}\{\theta_j\},$$

and the other coordinates θ_l^{++} for $l \neq 0$ are calculated as in (24) and (25). Similarly, $\boldsymbol{\theta}^{+-}, \boldsymbol{\theta}^{-+}$ can be defined with

$$\theta_j^{+-} = \theta_j^{+} = \theta_j - h_{jj} \text{sign}\{\theta_j\}.$$

When $\theta_0 = \theta_j = 0$ and $\theta_k \neq 0$, we use the gradient of $\boldsymbol{\theta}'$ for approximation, where $\boldsymbol{\theta}'$ is defined in (20). When $\theta_0 = \theta_j = \theta_k = 0$, we use the gradient of $\boldsymbol{\theta}''$ for approximation, where

$$\theta_0'' = 0, \quad \theta_j'' = \frac{1}{\sqrt{nT}}, \quad \theta_k'' = \frac{1}{\sqrt{nT}},$$

$$\theta_l'' = \theta_l \sqrt{1 - (\theta_j'')^2 - (\theta_k'')^2} \quad \text{for } l \neq 0, j, k.$$

Appendix B. Additional Details in Simulation Study

In this section, we provide the simulation details in Section 4.

B.1 MAD and Coverage Probability

In this section, we present the MAD and the coverage probability of the second important variable θ_2 and their average over the noise variables $\theta_{3:d}$ in Figures 3 and 4. The conclusion is similar to that for θ_1 . Note that the coverage probability of $\theta_{3:d}$ remains 100% for all

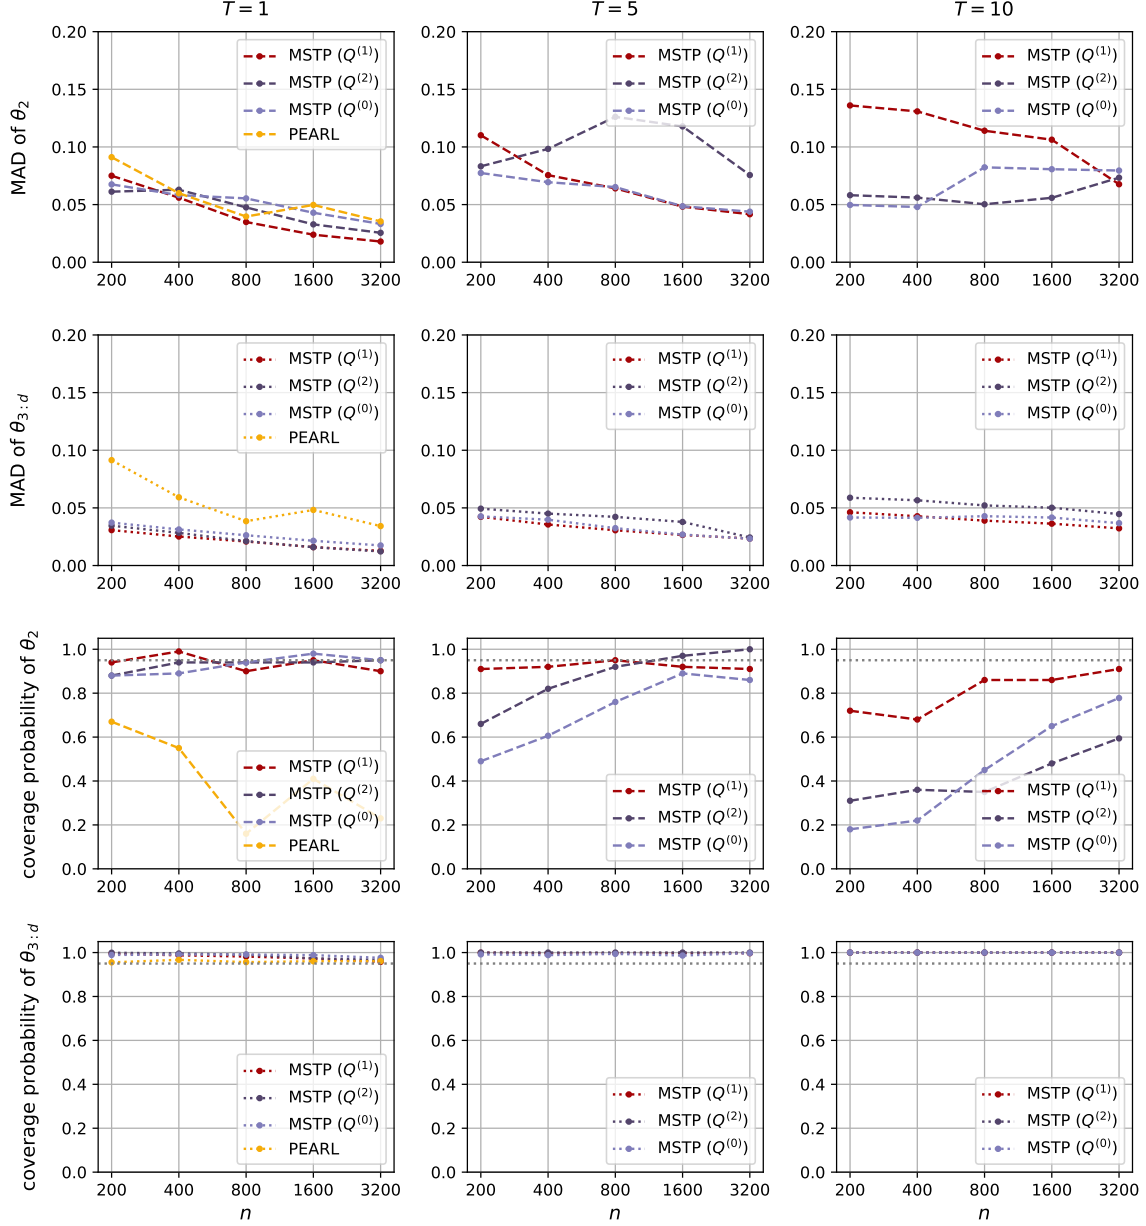


Figure 3: The MAD and coverage probability of θ_2 and their average over $\theta_{3:d}$ in Scenario 1. The columns correspond to different values of T , and the rows correspond to different metrics. The dotted line indicates the nominal coverage probability.

sample sizes when $T = 5$ and $T = 10$. This result is due to the bootstrap estimation procedure. Each bootstrap estimate $\tilde{\theta}_{b,j}$ of the one-step estimator is obtained by refitting the sparse estimator $\hat{\theta}_b$ on the bootstrap sample and conducting one-step estimation again. Since $\hat{\theta}_{b,j}$ for $j \in \{3 : d\}$ are estimated to be zero in most bootstrap samples, the one-step

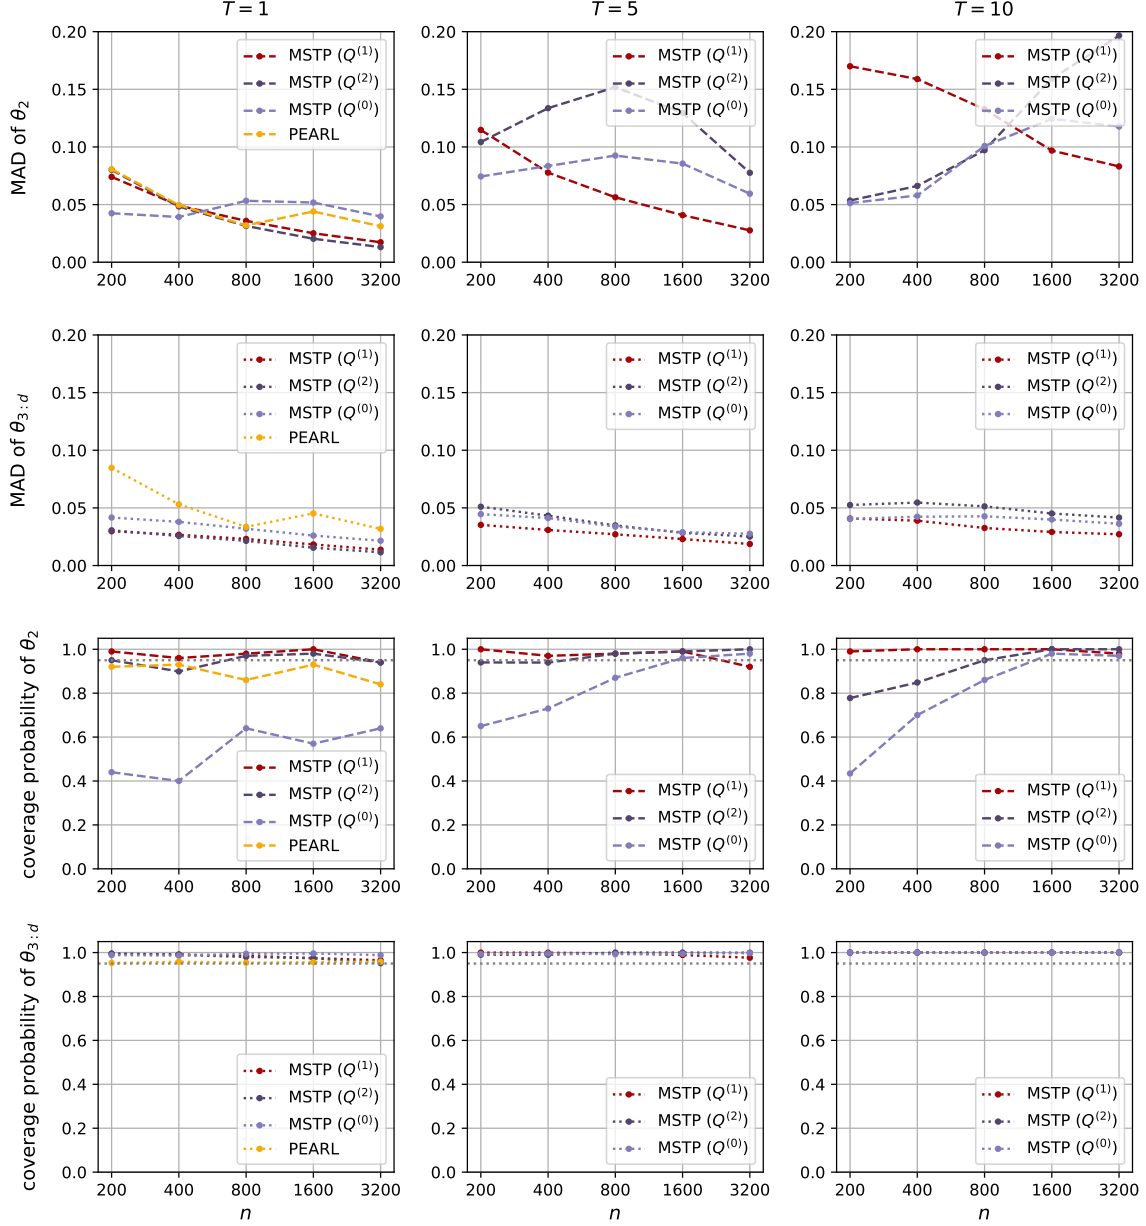


Figure 4: The MAD and coverage probability of θ_2 and their average over $\theta_{3:d}$ in Scenario 2. The columns correspond to different values of T , and the rows correspond to different metrics. The dotted line indicates the nominal coverage probability.

estimators cover the true value well enough. This is not an overestimation since the MAD of $\theta_{3:d}$ is much smaller than the MAD of θ_1 and θ_2 . For the one-step estimator of PEARL, after rescaling the CIs based on θ_1 , $\hat{\theta}_1$ achieves the nominal coverage probability. However, the CI of $\tilde{\theta}_2$ cannot cover θ^* well in Scenario 1. This is because the average ratio $\tilde{\theta}_2/\tilde{\theta}_1$ for PEARL changes from -0.226 when $n = 200$ to -0.675 when $n = 3200$.

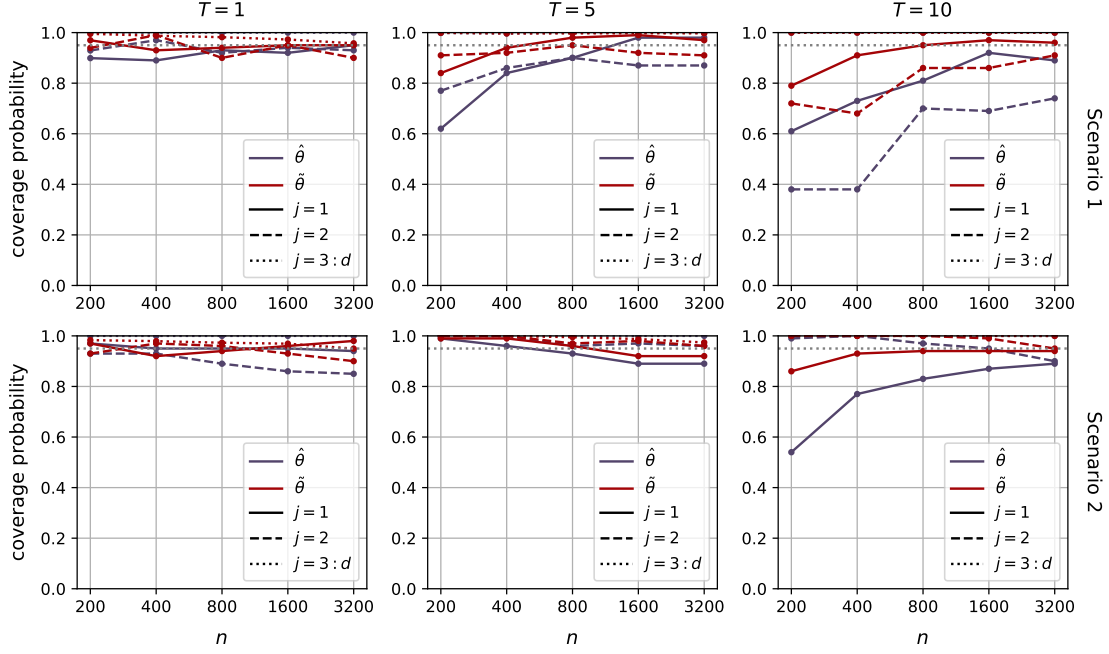


Figure 5: Comparison of the coverage probability of the CIs for the sparse estimator $\hat{\theta}$ and the one-step estimator $\tilde{\theta}$. The columns correspond to different values of T , and the rows correspond to different simulation scenarios.

B.2 Failure of the Sparse Estimator in Statistical Inference

In this section, we show that the sparse estimator $\hat{\theta}$ does not provide valid statistical inference results. Specifically, we use bootstrap to construct a CI for each $\hat{\theta}_j$, $j \in \{1 : d\}$, where the nuisance parameter is constructed using $Q^{(1)}$. Figure 5 provides the coverage probability of the CIs for $\hat{\theta}_j$ and $\tilde{\theta}_j$ for $j = 1, 2$, and the average coverage probability of $\hat{\theta}_{3:d}$ and $\tilde{\theta}_{3:d}$ over the noise variables. The results suggest that the CI based on $\hat{\theta}_j$ does not achieve the nominal coverage probability, especially when $T = 10$. This demonstrates the failure of the sparse estimator in statistical inference and the necessity of using the one-step estimator.

B.3 Validation of the Estimated MAD

In this section, we compare the estimated MAD and the empirical MAD of the MSTP estimated with $Q^{(1)}$. The estimated MAD refers to the MAD of the B bootstrap estimators $\tilde{\theta}_{b,j}$, averaged over W replications. The empirical MAD refers to the MAD of the W one-step estimators $\tilde{\theta}_j$. Figure 6 shows that the estimated MAD and the empirical MAD are approximately equal, verifying that the bootstrapped CI is a valid approximation for the CI of $\hat{\theta}_j$. Although the bootstrapped CI might be overly conservative when $T = 10$ and n is small in Scenario 1, the estimated MAD still recovers the empirical MAD as the sample size increases.

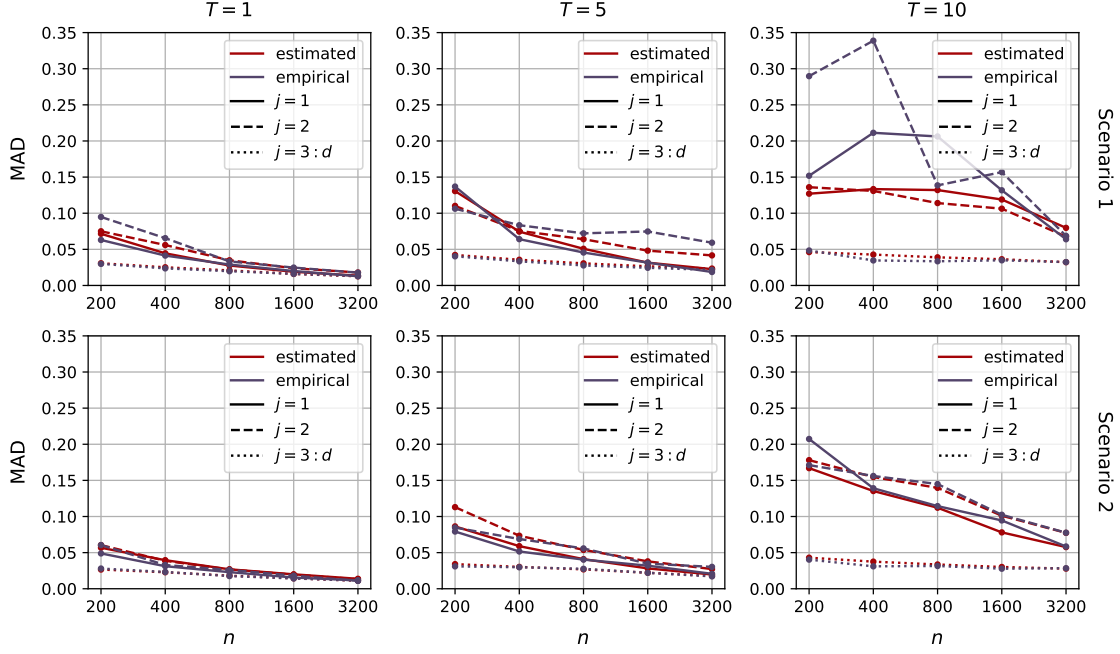


Figure 6: Comparison of the estimated and empirical MAD. The columns correspond to different values of T , and the rows correspond to different simulation scenarios.

B.4 True Value Function

Tables 2 and 3 provide the optimal parameter θ^* and the optimal average reward $\bar{V}^* = \max_{\theta} \bar{V}(\theta)$ in the simulation experiments in Section 4, where $\bar{V}(\theta) = \frac{1}{T} V(\theta)$. Figures 7 and 8 present the true average reward for all $\bar{V}(\theta)$. To find θ^* within the class Π , we estimate the average reward $\bar{V}(\theta)$ on an independent test set of size 200,000 for each $\theta = [\theta_0, \theta_1, \theta_2]$, where θ_1, θ_2 take values on the grids inside the unit ball and $\theta_0 = \pm\sqrt{1 - \theta_1^2 - \theta_2^2}$. The figures empirically verify that the value function is generally non-concave over the whole parameter space $\|\theta\|_2 \leq 1$, but it is locally concave around the true optimal parameter θ^* .

T	1	5	10
θ_0^*	-0.81	-0.38	-0.22
θ_1^*	0.53	0.67	0.61
θ_2^*	-0.26	-0.64	-0.76
V^*	2.200	3.455	5.791

Table 2: The optimal policy parameter θ^* and optimal average reward \bar{V}^* in Scenario 1.




Figure 7: The true value function in Scenario 1. In the first row, $\theta = [-\sqrt{1 - \theta_1^2 - \theta_2^2}, \theta_1, \theta_2]$; in the second row, $\theta = [\sqrt{1 - \theta_1^2 - \theta_2^2}, \theta_1, \theta_2]$.




Figure 8: The true value function in Scenario 2. In the first row, $\theta = [-\sqrt{1 - \theta_1^2 - \theta_2^2}, \theta_1, \theta_2]$; in the second row, $\theta = [\sqrt{1 - \theta_1^2 - \theta_2^2}, \theta_1, \theta_2]$.

T	1	5	10
θ_0^*	-0.84	-0.54	-0.33
θ_1^*	0.44	0.73	0.83
θ_2^*	0.31	0.42	0.45
V^*	0.794	1.599	2.978

Table 3: The optimal policy parameter θ^* and optimal average reward \bar{V}^* in Scenario 2.

Appendix C. Additional Details in Real Data Analysis

In this section, we present the details of the features, actions, and rewards. The latest OhioT1DM dataset contains two cohorts: a 2018 cohort and a 2020 cohort. The two cohorts used different fitness bands, which provide access to different physiological data. Since the basis heart rate, basis air temperature, basis steps, and sleep quality data are not available for the 2020 cohort, we focus on the 2018 cohort for analysis.

There are 17 relevant features recorded in the original dataset. We use the following 12 features to construct the MSTP. For a detailed description of the OhioT1DM dataset, please refer to Marling and Bunesu (2020).

- *Glucose level.* The CGM data, recorded every 5 minutes.
- *Basal.* The rate of continuous delivery of basal insulin. This rate remains constant for several hours until a new rate is set. The temporary basal variable is incorporated into this variable by superseding the normal basal rate.
- *Meal.* The self-reported carbohydrate estimate for meals. To represent the long-term effect of a meal, the carbohydrate estimate is passed along the next hour but discounted at a rate of 0.9. Specifically, if a patient has consumed a meal between the $(t - 1)$ th and the t th decision time, the feature is defined as $X_{\text{meal},s} = 0.9^{s-t} X_{\text{meal},t}$ for $s \in \{t : (t + 11)\}$.
- *Sleep.* The sleep quality at the decision time. If a patient is awake at this time, the variable takes a value of zero. This variable uses only the sleep quality detected by the fitness band to ensure access to real-time data.
- *Work.* The physical exertion during work time.
- *Hypoglycemic episode.* Whether a hypoglycemic event occurred before the decision time.
- *Exercise.* The physical exertion during exercise time.
- *Heart rate.* The heart rate, aggregated every 5 minutes.
- *GSR.* The galvanic skin response (GSR), aggregated every 5 minutes.
- *Skin temperature.* The skin temperature, aggregated every 5 minutes.
- *Air temperature.* The air temperature, aggregated every 5 minutes.

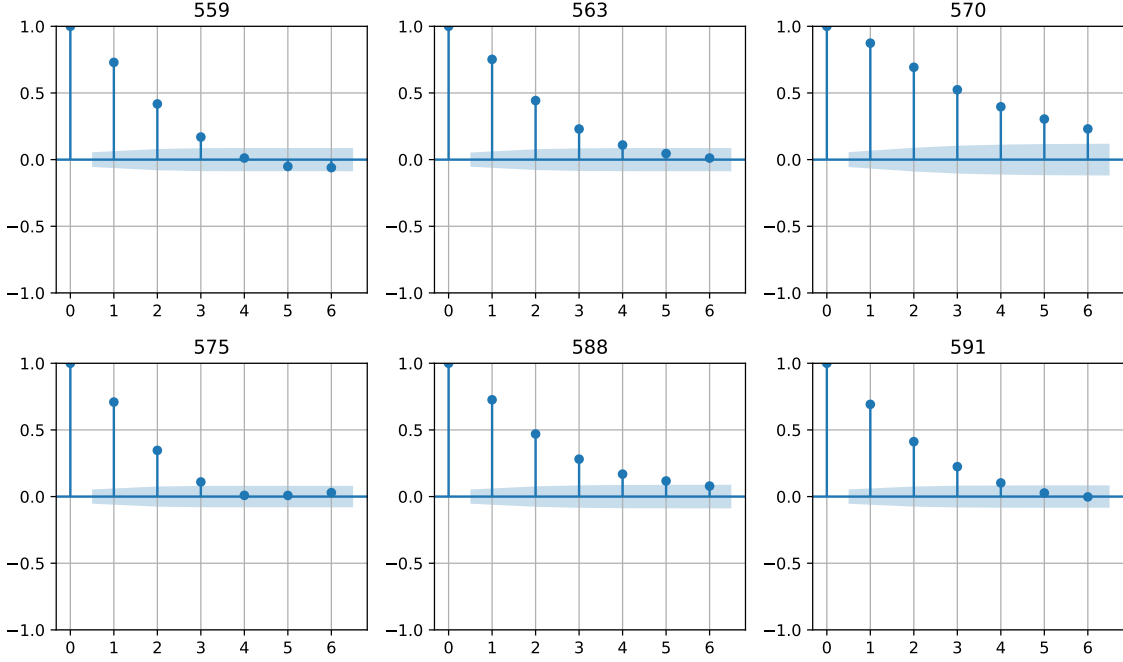


Figure 9: The autocorrelation of the initial glucose level for each hour across the 6 users in the 2018 cohort.

- *Steps*. The step count, aggregated every 5 minutes.

Although some features in the OhioT1DM dataset, including work time and its physical exertion, hypoglycemic episodes, and exercise time and duration, are self-reported after the event, a functional fitness band would be able to assess these features in real time. Hence, we still use these variables in estimating the MSTP, assuming that real-time data will be available to apply the estimated policy in future studies.

The self-reported sleep quality is omitted from the feature set since the automatic measure is already provided by the fitness band. The variable *finger stick* is omitted since it represents blood glucose values measured by self-monitoring and conveys similar information to the CGM data. The variable *stressor*, which represents the time of self-reported stress, is omitted because there were only 5 self-reported stress events across all 6 users. The variable *illness* is omitted because illness typically occurs over a period of time, but the ending time of each illness is missing. Additionally, there were only 11 self-reported illness events for the 6 users in total.

We define the treatment A_t based on whether bolus insulin is delivered or not, regardless of the dosage. Let j be the coordinate of the glucose level in the feature vector, and let $X_{t,j}$ represent the glucose level at time t . The reward is defined as the Index of Glycemic Control,

$$R_t = \begin{cases} -\frac{1}{30}(80 - X_{t+1,j})^2, & X_{t+1,j} < 80, \\ 0, & 80 \leq X_{t+1,j} \leq 140, \\ -\frac{1}{30}(X_{t+1,j} - 140)^{1.35}, & X_{t+1,j} \geq 140, \end{cases}$$

following Shi et al. (2020b).

Since we assume each trajectory is independent in theory, we separate the 8-week data of each user into multiple episodes. Figure 9 plots the autocorrelation function of the initial glucose level of each hour. We observe that the autocorrelation decreases to zero after 4 hours for most users. Therefore, we use the first hour of data from each 4-hour block for the real data analysis, while excluding the remaining 3 hours of data. In this way, temporal dependence across treatment stages is allowed within each trajectory, but independence is generally preserved between trajectories.

Appendix D. Theoretical Proofs

D.1 Gradient and Hessian Matrix of the Loss Function

In this section, we provide the theoretical gradient and Hessian matrix of the loss function. By the fact that any estimated $\boldsymbol{\theta}$ will satisfy the equation $\|\boldsymbol{\theta}\|_2 = 1$ to approximate a deterministic policy, we can rewrite the intercept as $\theta_0 = \pm\sqrt{1 - \theta_1^2 - \dots - \theta_d^2}$. When θ_0 is nonnegative, the gradient and Hessian matrix of the loss function are

$$\begin{aligned} \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}) = & -\frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \left\{ \frac{\nabla_{\boldsymbol{\xi}} [\prod_{k=1}^t \pi^{\boldsymbol{\theta}}(A_{i,k} | \mathbf{X}_{i,k})]}{\prod_{k=1}^t \mu_k(A_{i,k} | \mathbf{H}_{i,k})} [R_{i,t} - \hat{Q}_t(\mathbf{X}_{i,t}, A_{i,t})] \right. \\ & + \frac{\nabla_{\boldsymbol{\xi}} [\prod_{k=1}^{t-1} \pi^{\boldsymbol{\theta}}(A_{i,k} | \mathbf{X}_{i,k})]}{\prod_{k=1}^{t-1} \mu_k(A_{i,k} | \mathbf{H}_{i,k})} \sum_{a \in \mathcal{A}} \pi^{\boldsymbol{\theta}}(a | \mathbf{X}_{i,t}) \hat{Q}_t(\mathbf{X}_{i,t}, a), \\ & \left. + \frac{\prod_{k=1}^{t-1} \pi^{\boldsymbol{\theta}}(A_{i,k} | \mathbf{X}_{i,k})}{\prod_{k=1}^{t-1} \mu_k(A_{i,k} | \mathbf{H}_{i,k})} \sum_{a \in \mathcal{A}} \nabla_{\boldsymbol{\xi}} \pi^{\boldsymbol{\theta}}(a | \mathbf{X}_{i,t}) \hat{Q}_t(\mathbf{X}_{i,t}, a) \right\}, \end{aligned} \quad (26)$$

$$\begin{aligned} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}) = & -\frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \left\{ \frac{\nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 [\prod_{k=1}^t \pi^{\boldsymbol{\theta}}(A_{i,k} | \mathbf{X}_{i,k})]}{\prod_{k=1}^t \mu_k(A_{i,k} | \mathbf{H}_{i,k})} [R_{i,t} - \hat{Q}_t(\mathbf{X}_{i,t}, A_{i,t})] \right. \\ & + \frac{\nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 [\prod_{k=1}^{t-1} \pi^{\boldsymbol{\theta}}(A_{i,k} | \mathbf{X}_{i,k})]}{\prod_{k=1}^{t-1} \mu_k(A_{i,k} | \mathbf{H}_{i,k})} \sum_{a \in \mathcal{A}} \pi^{\boldsymbol{\theta}}(a | \mathbf{X}_{i,t}) \hat{Q}_t(\mathbf{X}_{i,t}, a), \\ & + 2 \frac{\nabla_{\boldsymbol{\xi}} [\prod_{k=1}^{t-1} \pi^{\boldsymbol{\theta}}(A_{i,k} | \mathbf{X}_{i,k})]}{\prod_{k=1}^{t-1} \mu_k(A_{i,k} | \mathbf{H}_{i,k})} \sum_{a \in \mathcal{A}} \nabla_{\boldsymbol{\xi}} \pi^{\boldsymbol{\theta}}(a | \mathbf{X}_{i,t}) \hat{Q}_t(\mathbf{X}_{i,t}, a), \\ & \left. + \frac{\prod_{k=1}^{t-1} \pi^{\boldsymbol{\theta}}(A_{i,k} | \mathbf{X}_{i,k})}{\prod_{k=1}^{t-1} \mu_k(A_{i,k} | \mathbf{H}_{i,k})} \sum_{a \in \mathcal{A}} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \pi^{\boldsymbol{\theta}}(a | \mathbf{X}_{i,t}) \hat{Q}_t(\mathbf{X}_{i,t}, a) \right\}, \end{aligned} \quad (27)$$

where

$$\begin{aligned}\nabla_{\boldsymbol{\xi}} \left[\prod_{k=1}^t \pi^{\boldsymbol{\theta}}(a_k | \mathbf{x}_k) \right] &= \left[\prod_{k=1}^t \pi^{\boldsymbol{\theta}}(a_k | \mathbf{x}_k) \right] \left[\sum_{k=1}^t \frac{a_k(\mathbf{x}_k - \boldsymbol{\xi}/\theta_0)/\tau}{1 + e^{a_k \mathbf{x}_k^T \boldsymbol{\theta}/\tau}} \right], \\ \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \left[\prod_{k=1}^t \pi^{\boldsymbol{\theta}}(a_k | \mathbf{x}_k) \right] &= \left[\prod_{k=1}^t \pi^{\boldsymbol{\theta}}(a_k | \mathbf{x}_k) \right] \left\{ \left[\sum_{k=1}^t \frac{a_k(\mathbf{x}_k - \boldsymbol{\xi}/\theta_0)/\tau}{1 + e^{a_k \mathbf{x}_k^T \boldsymbol{\theta}/\tau}} \right] \left[\sum_{k=1}^t \frac{a_k(\mathbf{x}_k - \boldsymbol{\xi}/\theta_0)/\tau}{1 + e^{a_k \mathbf{x}_k^T \boldsymbol{\theta}/\tau}} \right]^T \right. \\ &\quad \left. - \sum_{k=1}^t \frac{(1 + e^{a_k \mathbf{x}_k^T \boldsymbol{\theta}/\tau}) a_k / (\tau \theta_0) \mathbf{I} + e^{a_k \mathbf{x}_k^T \boldsymbol{\theta}/\tau} [a_k(\mathbf{x}_k - \boldsymbol{\xi}/\theta_0)/\tau] [a_k(\mathbf{x}_k - \boldsymbol{\xi}/\theta_0)/\tau]^T}{(1 + e^{a_k \mathbf{x}_k^T \boldsymbol{\theta}/\tau})^2} \right\}.\end{aligned}$$

Similarly, we can find the gradient and Hessian matrix when θ_0 is negative. Note that the nuisance parameter $\hat{\boldsymbol{\eta}} = \hat{Q}_{1:T}$ is estimated using the initial estimator $\hat{\boldsymbol{\theta}}$ and is then fixed in when minimizing the loss function. Therefore, the gradient and Hessian matrix with respect to $\boldsymbol{\xi}$ does not involve the gradient of $\hat{\boldsymbol{\eta}}$ with respect to $\boldsymbol{\xi}$.

D.2 Proof of Theorem 2

Since R_t is bounded for all t by Assumption 4, we know that $\bar{\boldsymbol{\eta}}$ is bounded by r with probability one. For a real valued function $f : \mathcal{D} \mapsto \mathbb{R}$, write the empirical process as

$$\mathbb{G}_n = \frac{1}{\sqrt{n}} \sum_{i=1}^n [f(\mathbf{D}_i) - \mathbb{E}f(\mathbf{D}_i)].$$

Denote $[\cdot]_j$ as the j th dimension of a vector and $[\cdot]_j$ the j th row of a matrix.

For $\mathcal{H} := \{\boldsymbol{\eta} - \bar{\boldsymbol{\eta}} : \boldsymbol{\eta} \in \mathcal{H}\}$ define the pathwise derivative of the nuisance parameter $D_q : \mathcal{H} \mapsto \mathbb{R}^d$,

$$D_q[\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}] := \nabla_q \mathbb{E} [\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}} + q(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}))], \quad \boldsymbol{\eta} \in \mathcal{H},$$

for all $q \in [0, 1]$. This derivative exists by our construction of ℓ . We will use the Neyman orthogonality to decorrelate the high-dimensional nuisance parameters. The definition is taken from Chernozhukov et al. (2018a).

Definition 3 (Neyman orthogonality) *The score $\nabla_{\boldsymbol{\xi}} \ell$ obeys the orthogonality condition at $(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})$ with respect to the nuisance realization set $\mathcal{H}_n \in \mathcal{H}$ if*

$$\mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})] = 0$$

and the pathwise derivative map $D_q[\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}]$ exists for all $q \in [0, 1]$ and $\boldsymbol{\eta} \in \mathcal{H}_n$ and vanishes at $q = 0$; namely,

$$D_0[\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}] = \mathbf{0}, \quad \text{for all } \boldsymbol{\eta} \in \mathcal{H}_n.$$

Lemma 4 *Under Assumptions 1-4, the gradient $\nabla_{\boldsymbol{\xi}} \ell$ satisfies the Neyman orthogonality. In addition, the following results hold for the nuisance parameters:*

$$\sup_{\boldsymbol{\eta} \in \mathcal{H}_n, q \in (0,1)} \left\| \nabla_{qq}^2 \mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}} + q(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}))] \right\|_{\infty} = 0, \quad (28)$$

$$\sup_{\boldsymbol{\eta} \in \mathcal{H}_n, q \in (0,1)} \left\| \nabla_q \mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}} + q(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}))] \right\|_{\infty} = 0. \quad (29)$$

Proof Since $\mathbb{E}[\ell(\boldsymbol{\theta}, \boldsymbol{\eta})] = -V(\boldsymbol{\theta})$ for any $\boldsymbol{\eta}$ and $\nabla_{\boldsymbol{\xi}} V(\boldsymbol{\theta}^*) = \mathbf{0}$ by the definition of $\boldsymbol{\theta}^*$, we have $\mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta})] = \mathbf{0}$ for any $\boldsymbol{\eta}$, when the regularity conditions are satisfied. Therefore,

$$D_0[\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}] = \nabla_{\boldsymbol{Q}} \mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta})]_{|\boldsymbol{\eta}=\bar{\boldsymbol{\eta}}}(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}) = \mathbf{0}$$

since $\nabla_{\boldsymbol{Q}} \mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta})] = \mathbf{0}$.

Similarly, for each dimension j we also have $\nabla_{\boldsymbol{Q}\boldsymbol{Q}}^2 \mathbb{E}[\nabla_{\theta_j} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta})] = \mathbf{0}$ since $\mathbb{E}[\nabla_{\theta_j} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta})]$ is a constant zero. Consequently,

$$\begin{aligned} & [\nabla_{qq}^2 \mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}} + q(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}))]]_j \\ &= \nabla_{qq}^2 \mathbb{E}[\nabla_{\theta_j} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}} + q(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}))] \\ &= (\boldsymbol{\eta} - \bar{\boldsymbol{\eta}})^T \nabla_{\boldsymbol{Q}\boldsymbol{Q}}^2 \mathbb{E}[\nabla_{\theta_j} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}} + q(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}))](\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}) \\ &= 0 \end{aligned}$$

and (28) follows. Since $\mathbb{E}[\nabla_{\boldsymbol{\xi}\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta})] = \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}} V(\boldsymbol{\theta}^*)$ for any $\boldsymbol{\eta}$, we have that

$$\nabla_q \mathbb{E}[\nabla_{\boldsymbol{\xi}\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}} + q(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}))] = \nabla_q \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}} V(\boldsymbol{\theta}^*) = 0$$

and thus (29) follows. ■

Lemma 5 *Under Assumptions 1-5, we have*

$$\sup_{\boldsymbol{\eta} \in \mathcal{H}_n} \|\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})\|_{\infty} = o_{\mathbb{P}}(1/\sqrt{n}), \quad (30)$$

$$\sup_{\boldsymbol{\eta} \in \mathcal{H}_n} \|\boldsymbol{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \boldsymbol{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})\|_{\infty} = o_{\mathbb{P}}(\sqrt{\log d/n}). \quad (31)$$

Proof To show (30), first note that

$$\|\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})\|_{\infty} \leq \frac{1}{\sqrt{n}} \|\mathbb{G}_n[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]\|_{\infty} \quad (32)$$

$$+ \|\mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]\|_{\infty} \quad (33)$$

for any $\boldsymbol{\eta} \in \mathcal{H}_n$. We will bound (33) using Neyman orthogonality and bound (32) using the results of empirical process.

To bound (33), define

$$\boldsymbol{h}_{\boldsymbol{\eta}}(q) := \mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}} + q(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}})) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]$$

for $q \in [0, 1)$. By Taylor's expansion, there exists $\tilde{q} \in (0, 1)$ such that

$$\boldsymbol{h}_{\boldsymbol{\eta}}(1) = \boldsymbol{h}_{\boldsymbol{\eta}}(0) + \nabla_q \boldsymbol{h}_{\boldsymbol{\eta}}(0) + \frac{1}{2} \nabla_{qq}^2 \boldsymbol{h}_{\boldsymbol{\eta}}(\tilde{q}).$$

We have $\boldsymbol{h}_{\boldsymbol{\eta}}(0) = \mathbf{0}$ by definition and $\nabla_q \boldsymbol{h}_{\boldsymbol{\eta}}(0) = \mathbf{0}$ by Lemma 4. In addition, the second derivative satisfies $\|\nabla_{qq}^2 \boldsymbol{h}_{\boldsymbol{\eta}}(\tilde{q})\|_{\infty} = 0$ by Lemma 4. Therefore, $\|\boldsymbol{h}_{\boldsymbol{\eta}}(1)\|_{\infty} = 0$ for any $\boldsymbol{\eta} \in \mathcal{H}_n$ and thus

$$\sup_{\boldsymbol{\eta} \in \mathcal{H}_n} \|\mathbb{E}[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]\|_{\infty} = o(1/\sqrt{n}).$$

To bound (32), define

$$\mathbf{g}_{\boldsymbol{\eta}_1, \boldsymbol{\eta}_2}(q) = \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}_1 + q(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1)) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}_1)$$

for any $\boldsymbol{\eta}_1, \boldsymbol{\eta}_2 \in \mathcal{H}_n$. Then we have

$$\mathbf{g}_{\boldsymbol{\eta}_1, \boldsymbol{\eta}_2}(1) = \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}_2) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}_1).$$

By Taylor's expansion, there exists $\tilde{q} \in (0, 1)$ such that

$$\mathbf{g}_{\boldsymbol{\eta}_1, \boldsymbol{\eta}_2}(1) = \mathbf{g}_{\boldsymbol{\eta}_1, \boldsymbol{\eta}_2}(0) + \nabla_q \mathbf{g}_{\boldsymbol{\eta}_1, \boldsymbol{\eta}_2}(\tilde{q}) = \nabla_{\boldsymbol{\theta}, \boldsymbol{\eta}}^2 \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}_1 + \tilde{q}(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1))(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1)$$

since $\mathbf{g}_{\boldsymbol{\eta}_1, \boldsymbol{\eta}_2}(0) = 0$ by definition. Hence by Cauchy-Schwartz inequality we have

$$[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}_2) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}_1)]_j \leq \|[\nabla_{\boldsymbol{\theta}, \boldsymbol{\eta}}^2 \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}_1 + \tilde{q}(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1))]_j\|_2 \|\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1\|_2 \quad (34)$$

for the j th dimension, which implies that the functions $[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta})]_j$ and $[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]_j$ are Lipschitz in the parameter $\boldsymbol{\eta}$. Note that with Assumption 4 and the boundedness of \hat{Q}_t , we have

$$\|\nabla_{\boldsymbol{\theta}, \boldsymbol{\eta}}^2 \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}_1 + \tilde{q}(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1))\|_2 \leq C,$$

where C is a constant.

Therefore, the bracketing number $N_{[]}(\epsilon, \mathcal{G}_{n,j}, L_2(\mathbb{P}))$ of the function set

$$\mathcal{G}_{n,j} := \{[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]_j : \boldsymbol{\eta} \in \mathcal{H}_n\}$$

is upper bounded by the covering number $N(\epsilon/(2C), \mathcal{H}_n, \|\cdot\|_2)$ of the nuisance parameter set \mathcal{H}_n (van der Vaart and Wellner, 1996, Theorem 2.7.11). Since $\|\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}\|_2 = o_{\mathbb{P}}(1)$, the covering number of the nuisance set $N(\epsilon/(2C), \mathcal{H}_n, \|\cdot\|_2)$ is finite. Let

$$J_{[]}(\delta, \mathcal{G}_{n,j}, \|\cdot\|) := \int_0^\delta \sqrt{1 + \log N_{[]}(\epsilon, \mathcal{G}_{n,j}, \|\cdot\|)} d\epsilon$$

be the bracketing integral. The minimum envelop function of class $\mathcal{G}_{n,j}$ is defined as $G_{n,j}(x) := \sup_{g \in \mathcal{G}_{n,j}} |g(x)|$. Suppose $\mathcal{G}_{n,j}$ is covered by the brackets $[l_1, u_1], \dots, [l_{N_\epsilon}, u_{N_\epsilon}]$, where $N_\epsilon := N_{[]}(\epsilon, \mathcal{G}_{n,j}, L_2(\mathbb{P}))$ is the bracketing number for any $\epsilon > 0$. Then we can write the minimum envelop function as

$$G_{n,j}(x) = \max_{j \in \{1:n\}_\epsilon} \{|l_{N_\epsilon}(x)|, |u_{N_\epsilon}(x)|\}.$$

The $L_{\mathbb{P},2}$ norm of $G_{n,j}(x)$ is then

$$\|G_{n,j}\|_{\mathbb{P},2}^2 = \mathbb{E} G_{n,j}^2(\mathbf{D}) \leq \mathbb{E} \sum_{j=1}^{N_\epsilon} (l_j^2(\mathbf{D}) + u_j^2(\mathbf{D})).$$

Note that

$$\begin{aligned} \mathbb{E} l_j^2(\mathbf{D}) &= \mathbb{E} [g(\mathbf{D}) + (l_j(\mathbf{D}) - g(\mathbf{D}))]^2 \\ &\leq 2[\mathbb{E} g^2(\mathbf{D}) + \mathbb{E} (l_j(\mathbf{D}) - g(\mathbf{D}))^2] \\ &\leq 2[\mathbb{E} g^2(\mathbf{D}) + \epsilon^2] \end{aligned}$$

for some $g \in \mathcal{G}_{n,j}$ contained in the j th bracket for any $j \in \{1 : N_\epsilon\}$. According to (34),

$$\mathbb{E}g^2(\mathbf{D}) \leq C^2 \|\boldsymbol{\eta} - \bar{\boldsymbol{\eta}}\|_{\mathbb{P},2} \leq \delta_n$$

by Assumption 5 and thus

$$\mathbb{E}l_j^2(\mathbf{D}) \leq 2[\epsilon^2 + \delta_n].$$

The same holds for upper brackets u_j . Since

$$N_\epsilon \leq N \left(\frac{\epsilon}{2C}, \mathcal{H}_n, \|\cdot\|_2 \right) \leq \delta_n \frac{2C}{\epsilon},$$

we get that

$$\|G_{n,j}\|_{\mathbb{P},2}^2 \leq 4[\epsilon^2 + \delta_n]N_\epsilon \leq 4[\epsilon^2 + \delta_n] \cdot 2C\delta_n \frac{1}{\epsilon}.$$

So $\|G_{n,j}\|_{\mathbb{P},2}^2 \rightarrow 0$ when $n \rightarrow \infty$. Since

$$\mathbb{E}^* \sup_{g \in \mathcal{G}_{n,j}} \mathbb{G}_n(g) \lesssim J_{[]}(\|G_{n,j}\|_{\mathbb{P},2}, \mathcal{G}_{n,j}, L_2(\mathbb{P})) \quad (35)$$

by van der Vaart and Wellner (1996, Theorem 2.14.2), the left-hand side of (35) is in the order of $o(1)$. Finally, by Markov's inequality,

$$\sup_{\boldsymbol{\eta} \in \mathcal{H}_n} \frac{1}{\sqrt{n}} \mathbb{G}_n[\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]_j = o_{\mathbb{P}}(1/\sqrt{n}).$$

The bound on (32) follows by taking the maximum over all the dimensions.

Combining the upper bounds of (33) and (32) we can show (30).

For (31), note that

$$\|\mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})\|_{\infty} \quad (36)$$

$$\leq \frac{1}{\sqrt{n}} \|\mathbb{G}_n[\mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]\|_{\infty} \quad (37)$$

$$+ \|\mathbb{E}[\mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \boldsymbol{\eta}) - \mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]\|_{\infty} \quad (38)$$

for any $\boldsymbol{\eta} \in \mathcal{H}_n$.

To bound (38), define

$$\tilde{\mathbf{h}}_{\boldsymbol{\eta}}(q) := \mathbb{E}[\mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}} + q(\boldsymbol{\eta} - \bar{\boldsymbol{\eta}})) - \mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]$$

for $q \in [0, 1)$. By Taylor's expansion, there exists $\tilde{q} \in (0, 1)$ such that

$$\tilde{\mathbf{h}}_{\boldsymbol{\eta}}(1) = \tilde{\mathbf{h}}_{\boldsymbol{\eta}}(0) + \nabla_q \tilde{\mathbf{h}}_{\boldsymbol{\eta}}(\tilde{q}).$$

We have $\tilde{\mathbf{h}}_{\boldsymbol{\eta}}(0) = \mathbf{0}$ by definition and $\nabla_q \tilde{\mathbf{h}}_{\boldsymbol{\eta}}(\tilde{q}) = o(\sqrt{\log d/n})$ by Lemma 4.

Using similar arguments as that for (32), we can conclude that (37) is in the order of $o_{\mathbb{P}}(1/\sqrt{n})$. Combining the upper bounds of (38) and (37) gives the results. \blacksquare

Lemma 6 (Concentration of the gradient and Hessian) *Under Assumptions 1-5, we have*

$$\|\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbb{E}\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})\|_{\infty} = \mathcal{O}_{\mathbb{P}}(\sqrt{\log d/n}), \quad (39)$$

$$\|\mathbf{v}_j^{*T} \nabla_{\xi\xi}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbb{E}(\mathbf{v}_j^{*T} \nabla_{\xi\xi}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}}))\|_{\infty} = \mathcal{O}_{\mathbb{P}}(\sqrt{\log d/n}). \quad (40)$$

Proof Since $\mathbb{P}(\hat{\boldsymbol{\eta}} \notin \mathcal{H}_n) \leq \Delta_n$ and Δ_n converges to zero, we can focus on the event when $\boldsymbol{\eta} \in \mathcal{H}_n$. To prove (39), note that

$$\begin{aligned} & \|\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbb{E}\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})\|_{\infty} \\ & \leq \|\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \nabla_{\xi}\ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})\|_{\infty} + \|\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}}) - \mathbb{E}\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})\|_{\infty}. \end{aligned}$$

By (30) in Lemma 5, we only need to show that

$$\|\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}}) - \mathbb{E}\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})\|_{\infty} = \mathcal{O}_{\mathbb{P}}(\sqrt{\log d/n}). \quad (41)$$

Write

$$\mathbf{h}(\mathbf{D}_{1:n}) := -\nabla_{\xi}\ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}}),$$

which is a d -dimensional real-valued function, and denote the j -th dimension of \mathbf{h} as h_j . According to (26), we have

$$\begin{aligned} \mathbf{h}(\mathbf{D}_{1:n}) = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T & \left\{ \nabla_{\xi} \rho_{i,1:t}^{\boldsymbol{\theta}^*, \bar{\boldsymbol{\mu}}} [R_{i,t} - \bar{\eta}_t(\mathbf{X}_{i,t}, A_{i,t})] + [\nabla_{\xi} \rho_{i,1:(t-1)}^{\boldsymbol{\theta}^*, \bar{\boldsymbol{\mu}}}] \bar{U}_t(\mathbf{X}_{i,t}) \right. \\ & \left. + \rho_{i,1:(t-1)}^{\boldsymbol{\theta}^*, \bar{\boldsymbol{\mu}}} [\nabla_{\xi} \bar{U}_t(\mathbf{X}_{i,t})] \right\}, \end{aligned}$$

where $\bar{U}_t(\mathbf{X}_{i,t}) = \sum_{a \in \mathcal{A}} \pi^{\boldsymbol{\theta}}(a, \mathbf{X}_{i,t}) \bar{\eta}_t(\mathbf{X}_{i,t}, a)$. When $\theta_0^* \geq 0$, since

$$\rho_{i,1:t}^{\boldsymbol{\theta}^*, \bar{\boldsymbol{\mu}}} \leq \frac{1}{p_0^t} \quad \text{and} \quad \left| \frac{A_{i,k}(X_{i,k,j} - \theta_j^*/\theta_0^*)/\tau}{1 + e^{A_{i,k} \mathbf{X}_{i,k}^T \boldsymbol{\theta}^*/\tau}} \right| \leq \frac{z - u^{-1}}{\tau},$$

we have that the j th dimension of $\nabla_{\xi} \rho_{i,1:t}^{\boldsymbol{\theta}^*, \bar{\boldsymbol{\mu}}}$ is upper bounded by $\frac{zt}{\tau p_0^t}$. Besides, we know that $|R_{i,t} - \bar{\eta}_t(\mathbf{X}_{i,t}, A_{i,t})| \leq 2r$ since $\bar{\eta}_t$ is bounded by r according to Assumption 4 and the definition of $\bar{\eta}_t$. Thus the j th dimension of the first part is changed by

$$\begin{aligned} & \frac{1}{n} \sum_{t=1}^T \left| [\nabla_{\xi} \rho_{i,1:t}^{\boldsymbol{\theta}^*, \bar{\boldsymbol{\mu}}}]_j [R_{i,t} - \bar{\eta}_t(\mathbf{X}_{i,t}, A_{i,t})] - [\nabla_{\xi} \rho_{i,1:t}^{\boldsymbol{\theta}^*, \bar{\boldsymbol{\mu}}}]_j [R'_{i,t} - \bar{\eta}_t(\mathbf{X}'_{i,t}, A'_{i,t})] \right| \\ & \leq \frac{2}{n} \sum_{t=1}^T \frac{2rt(z - u^{-1})}{\tau p_0^t} \end{aligned}$$

if the i th trajectory \mathbf{D}_i is changed into \mathbf{D}'_i . Similarly, since $||[\nabla_{\xi} \pi^{\boldsymbol{\theta}^*}(a|\mathbf{x})]_j| \leq \frac{z - u^{-1}}{\tau}$, we have

$$|[\nabla_{\xi} \rho_{i,1:(t-1)}^{\boldsymbol{\theta}^*, \bar{\boldsymbol{\mu}}}]_j \bar{U}_t(\mathbf{X}_{i,t})| \leq \frac{r(t-1)(z - u^{-1})}{\tau p_0^{t-1}}$$

and

$$|\rho_{i,1:(t-1)}^{\boldsymbol{\theta}^*, \bar{\boldsymbol{\mu}}} [\nabla_{\boldsymbol{\xi}} \bar{U}_t(\mathbf{X}_{i,t})]_j| \leq \frac{2r(z - u^{-1})}{\tau p_0}.$$

Therefore, the upper bound on the change of h_j when changing \mathbf{D}_i is $c_i := C/n$ for some constant $C > 0$ depending on r, z, p_0, τ, T . We have similar results when $\theta_0^* < 0$. By McDiarmid's inequality, we get that

$$\mathbb{P}(|h_j(\mathbf{D}_{1:n}) - \mathbb{E}h_j(\mathbf{D}_{1:n})| \geq \epsilon) \leq 2 \exp \left\{ -\frac{2\epsilon^2}{\sum_{i=1}^n c_i^2} \right\} \leq 2 \exp \{-C'n\epsilon^2\}$$

for some constant $C' > 0$. Now the union bound inequality yields

$$\mathbb{P}(\|\mathbf{h}(\mathbf{D}_{1:n}) - \mathbb{E}\mathbf{h}(\mathbf{D}_{1:n})\|_{\infty} \geq \epsilon) \leq 2d \exp \{-C'n\epsilon^2\}.$$

With $\epsilon = C''\sqrt{\log d/n}$ for some $C'' > 0$, we have $\mathbf{h}(\mathbf{D}_{1:n}) = \mathcal{O}_{\mathbb{P}}(\sqrt{\log d/n})$ and equation (41) follows.

Similarly, by (31) in Lemma 5, we only need to prove

$$\|\mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}}) - \mathbb{E}(\mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}}))\|_{\infty} = \mathcal{O}_{\mathbb{P}}(\sqrt{\log d/n})$$

for (40). Now $\mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})$ can be divided into 5 parts. When $\bar{\boldsymbol{\eta}}_t$ is upper bounded by the constant r , each dimension for each part is bounded by C''/n for some constant $C'' > 0$, since \mathbf{v}^* is a constant. The result follows from the same arguments as before. \blacksquare

Lemma 7 (Central limit theorem for the score function) *Under Assumptions 1-3 and 7, it holds that*

$$\sqrt{n} \mathbf{v}_j^{*T} \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \Rightarrow N(0, \sigma_j^*),$$

where $\sigma_j^* \geq C$ for some constant $C > 0$.

Proof Since $\mathbb{E} \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}}) = 0$, note that

$$\sqrt{n} \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) = \sqrt{n} [\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})] + \sqrt{n} [\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}}) - \mathbb{E} \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})]. \quad (42)$$

The equation (30) in Lemma 5 shows that the first difference on the right-hand side of (42) is in the order of $o_{\mathbb{P}}(1)$ when $\hat{\boldsymbol{\eta}} \in \mathcal{H}_n$. Besides, the probability of $\hat{\boldsymbol{\eta}} \notin \mathcal{H}_n$ converges to zero. For the second difference in (42), when $\boldsymbol{\Sigma}^*$ is finite by Assumption 7, the multivariate central limit theorem (Ferguson, 2017, Theorem 5) shows that

$$\sqrt{n} [\nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}}) - \mathbb{E} \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \bar{\boldsymbol{\eta}})] \Rightarrow N(0, \boldsymbol{\Sigma}^*).$$

The convergence follows since $N(0, \mathbf{v}_j^{*T} \boldsymbol{\Sigma}^* \mathbf{v}_j^*) + o_{\mathbb{P}}(1) = N(0, \mathbf{v}_j^{*T} \boldsymbol{\Sigma}^* \mathbf{v}_j^*)$. In addition, Assumption 7 guarantees that $\mathbf{v}_j^{*T} \boldsymbol{\Sigma}^* \mathbf{v}_j^* \geq C$, since \mathbf{v}_j^{*T} is nonzero at least in its first argument. \blacksquare

Lemma 8 *Under Assumptions 1-6, when $\lambda_{\theta} \simeq \sqrt{\log d/n}$, we have*

$$(\hat{\xi} - \xi^*)^T \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta})(\hat{\xi} - \xi^*) = \mathcal{O}_{\mathbb{P}}(s_{\xi} \log d/n). \quad (43)$$

Here, we present the proof for Lemmas 1 and 8 together.

Proof Denote $\hat{\Delta}_{\xi} := \hat{\xi} - \xi^*$. Let \mathcal{S}_{ξ} and \mathcal{S}_{w_j} be the support set of ξ^* and w_j^* respectively. According to Chernozhukov et al. (2018b, inequality (88)), we have

$$\begin{aligned} \langle \nabla_{\xi} \ell(\hat{\theta}, \hat{\eta}), \hat{\Delta}_{\xi} \rangle &\leq \lambda_{\xi} (\|\xi^*\|_1 - \|\hat{\xi}\|_1) = \lambda_{\xi} (\|\xi_{\mathcal{S}_{\xi}}^*\|_1 - \|\xi_{\mathcal{S}_{\xi}}^* + \hat{\Delta}_{\xi, \mathcal{S}_{\xi}}\|_1 - \|\hat{\Delta}_{\xi, \mathcal{S}_{\xi}^c}\|_1) \\ &\leq \lambda_{\xi} (\|\hat{\Delta}_{\xi, \mathcal{S}_{\xi}}\|_1 - \|\hat{\Delta}_{\xi, \mathcal{S}_{\xi}^c}\|_1). \end{aligned} \quad (44)$$

Define the empirical symmetric Bregman distance as

$$H(\xi, \xi^*, \eta) = \langle \nabla_{\xi} \ell(\theta, \eta) - \nabla_{\xi} \ell(\theta^*, \eta), \xi - \xi^* \rangle.$$

Since all variables are bounded, by Taylor's expansion on each dimension we have

$$\sup_{\xi} \|\nabla_{\xi\xi} \ell(\theta, \eta) - \nabla_{\xi\xi} \ell(\theta, \eta')\|_{\infty} \leq L \|\eta - \eta'\|_2$$

for any $\eta, \eta' \in \mathcal{H}_n$. Therefore, the empirical symmetric Bregman distance is Lipschitz in the nuisance parameters:

$$|H(\xi, \xi^*, \eta) - H(\xi, \xi^*, \eta')| \leq L \|\eta - \eta'\|_2 \|\xi - \xi^*\|_1^2 \quad (45)$$

for any $\eta, \eta' \in \mathcal{H}_n$ and for any ξ (Chernozhukov et al., 2018b, Lemma 3).

By Assumption 6, we know $\mathbb{E} \ell(\theta, \bar{\eta}) = -V(\xi)$ is κ -strongly convex at ξ^* . We will use Chernozhukov et al. (2018b, Lemma 2) to prove the restricted strong convexity:

$$H(\xi^* + \Delta_{\xi}, \xi^*, \bar{\eta}) \geq \kappa \|\Delta_{\xi}\|_2^2 - \tau_{n,\delta} \|\Delta_{\xi}\|_1^2 \quad (46)$$

holds with probability $1 - \delta$ for all $\Delta_{\xi} \in \mathbb{B}(R)$ and $\tau_{n,\delta} \simeq 1/(\delta\sqrt{n})$. To verify the condition, we need to show that

$$\sup_{\xi \in \{\xi^* + \Delta_{\xi} : \Delta_{\xi} \in \mathbb{B}(R)\}} \|\mathbf{v}_j^{*T} \nabla_{\xi\xi} \ell(\theta, \bar{\eta}) - \mathbf{v}_j^{*T} \mathbb{E} \nabla_{\xi\xi} \ell(\theta, \bar{\eta})\|_{\infty} \leq \tau_{n,\delta}$$

with probability $1 - \delta$. Similar as the bound for (32), for each dimension j , we can use Taylor's expansion to show that $\nabla_{\xi\xi} \ell(\theta, \bar{\eta})$ is Lipschitz in the parameter ξ with constant C . Consequently, the bracketing number $N_{[]}(\epsilon, \mathcal{G}_{n,j}, L_2(\mathbb{P}))$ of

$$\mathcal{G}_{n,j} := \{\mathbf{v}_j^{*T} \nabla_{\xi\xi} \ell(\theta^* + \Delta_{\xi}, \bar{\eta})\}_j : \Delta_{\xi} \in \mathbb{B}(R)\}$$

is upper bounded by the covering number $N(\epsilon/(2C), \mathbb{B}(R), \|\cdot\|_2)$. For a finite R , the set $\mathbb{B}(R)$ is bounded and $N(\epsilon, \mathbb{B}(R), \|\cdot\|_2) \simeq (1/\epsilon)^d$. Therefore,

$$\mathbb{E}^* \sup_{g \in \mathcal{G}_{n,j}} \mathbb{G}_n(g) \lesssim J_{[]}(\|G_{n,j}\|_{\mathbb{P},2}, \mathcal{G}_{n,j}, L_2(\mathbb{P})) \leq J_{[]}(\infty, \mathcal{G}_{n,j}, L_2(\mathbb{P})) < \infty$$

by van der Vaart and Wellner (1996, Theorem 2.14.2). By Markov's inequality and union bound inequality, we get that

$$\tau_{n,\delta} \simeq dJ_{\square}(\|G_{n,j}\|_{\mathbb{P},2}, \mathcal{G}_{n,j}, L_2(\mathbb{P})) / (\delta\sqrt{n}).$$

Combining the Lipschitz bound (45) and the restricted strong convexity bound (46), we have

$$\begin{aligned} H(\hat{\xi}, \xi^*, \hat{\eta}) &\geq H(\hat{\xi}, \xi^*, \bar{\eta}) - L\|\hat{\eta} - \bar{\eta}\|_2 \|\hat{\Delta}_{\xi}\|_1^2 \\ &\geq \kappa \|\hat{\Delta}_{\xi}\|_2^2 - (\tau_{n,\delta} + L\|\hat{\eta} - \bar{\eta}\|_2) \|\hat{\Delta}_{\xi}\|_1^2. \end{aligned}$$

On the other hand,

$$\begin{aligned} H(\hat{\xi}, \xi^*, \hat{\eta}) &\leq \left\langle \nabla_{\xi} \ell(\hat{\theta}, \hat{\eta}) - \nabla_{\xi} \ell(\theta^*, \hat{\eta}), \hat{\xi} - \xi^* \right\rangle \\ &\leq \lambda_{\xi} (\|\hat{\Delta}_{\xi, S_{\xi}}\|_1 - \|\hat{\Delta}_{\xi, S_{\xi}^c}\|_1) + \|\nabla_{\xi} \ell(\theta^*, \hat{\eta})\|_{\infty} \|\hat{\Delta}_{\xi}\|_1 \end{aligned}$$

according to (44). The assumption that $\lambda_{\xi} \simeq \sqrt{\log d/n}$ and Lemma 6 implies

$$\|\nabla_{\xi} \ell(\theta^*, \hat{\eta})\|_{\infty} \lesssim \lambda_{\xi}/2$$

with high probability. Conditioning on this event, the above two bounds yield

$$\begin{aligned} \kappa \|\hat{\Delta}_{\xi}\|_2^2 &\leq \frac{\lambda_{\xi}}{2} (3\|\hat{\Delta}_{\xi, S_{\xi}}\|_1 - \|\hat{\Delta}_{\xi, S_{\xi}^c}\|_1) + (\tau_{n,\delta} + L\|\hat{\eta} - \bar{\eta}\|_2) \|\hat{\Delta}_{\xi}\|_1^2 \\ &\leq \frac{3\lambda_{\xi}}{2} \sqrt{s_{\xi}} \|\hat{\Delta}_{\xi}\|_2 + (\tau_{n,\delta} + L\|\hat{\eta} - \bar{\eta}\|_2) s_{\xi} \|\hat{\Delta}_{\xi}\|_2^2. \end{aligned}$$

Since $\mathbb{P}(\eta \notin \mathcal{H}_n) \rightarrow 0$, we focus on the events when $\eta \in \mathcal{H}_n$. When n is large enough such that $(\tau_{n,\delta} + L\delta_n)s_{\xi} \leq \kappa/2$, we have

$$\|\hat{\Delta}_{\xi}\|_2 \leq \frac{3\lambda_{\xi}}{4} \sqrt{s_{\xi}} \simeq \sqrt{s_{\xi} \log d/n}.$$

Besides, since

$$\frac{\lambda_{\xi}}{2} (3\|\hat{\Delta}_{\xi, S_{\xi}}\|_1 - \|\hat{\Delta}_{\xi, S_{\xi}^c}\|_1) \geq \frac{\kappa}{2} \|\hat{\Delta}_{\xi}\|_2^2 \geq 0,$$

we also have $\|\hat{\Delta}_{\xi, S_{\xi}^c}\|_1 \leq 3\|\hat{\Delta}_{\xi, S_{\xi}}\|_1$ and thus

$$\|\hat{\Delta}_{\xi}\|_1 \leq 4\|\hat{\Delta}_{\xi, S_{\xi}}\|_1 \leq 4\sqrt{s_{\xi}} \|\hat{\Delta}_{\xi, S_{\xi}}\|_2 \lesssim s_{\xi} \sqrt{\log d/n}.$$

Now consider the randomness of $\|\nabla_{\xi} \ell(\theta^*, \hat{\eta})\|_{\infty}$ and we have $\|\hat{\Delta}_{\xi}\|_1 = \mathcal{O}_{\mathbb{P}}(s_{\xi} \sqrt{\log d/n})$.

To show (43), note that

$$\begin{aligned} &\left| H(\hat{\xi}, \xi^*, \hat{\eta}) - \hat{\Delta}_{\xi}^T \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) \hat{\Delta}_{\xi} \right| \\ &= \left| \hat{\Delta}_{\xi}^T \left[\nabla_{\xi\xi}^2 \ell(\tilde{q}\hat{\theta} + (1-\tilde{q})\theta^*, \hat{\eta}) - \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) \right] \hat{\Delta}_{\xi} \right| \\ &\leq \left\| \hat{\Delta}_{\xi}^T \left[\nabla_{\xi\xi}^2 \ell(\tilde{q}\hat{\theta} + (1-\tilde{q})\theta^*, \hat{\eta}) - \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) \right] \right\|_{\infty} \|\hat{\Delta}_{\xi}\|_1 \end{aligned}$$

by Taylor's expansion for some $\tilde{q} \in [0, 1]$ as in the proof for Lemma 5. Define

$$g_j(q') = \hat{\Delta}_\xi^T \left[\nabla_{\xi\theta_j}^2 \ell(q'[\tilde{q}\hat{\theta} + (1 - \tilde{q})\theta^*] + (1 - q')\theta^*, \hat{\eta}) \right]$$

for coordinate j . Then we have

$$g_j(1) = \hat{\Delta}_\xi^T \left[\nabla_{\xi\theta_j}^2 \ell(\tilde{q}\hat{\theta} + (1 - \tilde{q})\theta^*, \hat{\eta}) \right],$$

and

$$g_j(0) = \hat{\Delta}_\xi^T \left[\nabla_{\xi\theta_j}^2 \ell(\theta^*, \hat{\eta}) \right].$$

With mean value theorem again, there exists \tilde{q}' such that

$$g_j(1) = g_j(0) + \frac{d}{dq'} g_j(\tilde{q}').$$

i.e.,

$$\begin{aligned} & \hat{\Delta}_\xi^T \left[\nabla_{\xi\theta_j}^2 \ell(\tilde{q}\hat{\theta} + (1 - \tilde{q})\theta^*, \hat{\eta}) - \nabla_{\xi\theta_j}^2 \ell(\theta^*, \hat{\eta}) \right] \\ &= \hat{\Delta}_\xi^T \nabla_\xi \left[\nabla_{\xi\theta_j}^2 \ell(\tilde{q}'[\tilde{q}\hat{\theta} + (1 - \tilde{q})\theta^*] + (1 - \tilde{q}')\theta^*, \hat{\eta}) - \nabla_{\xi\theta_j}^2 \ell(\theta^*, \hat{\eta}) \right] \tilde{q} \hat{\Delta}_\xi \end{aligned} \quad (47)$$

for some $\tilde{q}' \in [0, 1]$. Since all variables are bounded, we have

$$\hat{\Delta}_\xi^T \nabla_\xi \left[\nabla_{\xi\theta_j}^2 \ell(\tilde{q}'[\tilde{q}\hat{\theta} + (1 - \tilde{q})\theta^*] + (1 - \tilde{q}')\theta^*, \hat{\eta}) - \nabla_{\xi\theta_j}^2 \ell(\theta^*, \hat{\eta}) \right] \hat{\Delta}_\xi \leq C \hat{\Delta}_\xi^T \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) \hat{\Delta}_\xi$$

for some constant C for all j . Hence,

$$\left| H(\hat{\xi}, \xi^*, \hat{\eta}) - \hat{\Delta}_\xi^T \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) \hat{\Delta}_\xi \right| \leq C \hat{\Delta}_\xi^T \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) \hat{\Delta}_\xi \left\| \hat{\Delta}_\xi \right\|_1 \quad (48)$$

Since $\left\| \hat{\Delta}_\xi \right\|_1 = \mathcal{O}_{\mathbb{P}}(s_\xi \sqrt{\log d/n})$,

$$\hat{\Delta}_\xi^T \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) \hat{\Delta}_\xi (1 - Cs_\xi \sqrt{\log d/n}) \lesssim H(\hat{\xi}, \xi^*, \hat{\eta}).$$

Finally, since

$$H(\hat{\xi}, \xi^*, \hat{\eta}) \leq \frac{\lambda_\xi}{2} (3 \left\| \hat{\Delta}_{\xi, S_\xi} \right\|_1 - \left\| \hat{\Delta}_{\xi, S_\xi^c} \right\|_1) \lesssim \lambda_\xi \left\| \hat{\Delta}_{\xi, S_\xi} \right\|_1 = \mathcal{O}_{\mathbb{P}}(s_\xi \log d/n),$$

we conclude that $\hat{\Delta}_\xi^T \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) \hat{\Delta}_\xi = \mathcal{O}_{\mathbb{P}}(s_\xi \log d/n)$. ■

Lemma 9 *Under Assumptions 1-5, when $\lambda_{w_j} \simeq \sqrt{\log d/n}$, we have*

$$\left\| \hat{w}_j - w_j^* \right\|_1 = \mathcal{O}_{\mathbb{P}}((s_\xi \vee s_{w_j}) \sqrt{\log d/n}), \quad (49)$$

$$(\hat{v}_j - v_j^*)^T \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) (\hat{v}_j - v_j^*) = \mathcal{O}_{\mathbb{P}}((s_\xi \vee s_{w_j}) \log d/n), \quad (50)$$

$$(\hat{v}_j - v_j^*)^T \nabla_{\xi\xi}^2 \ell(\hat{\theta}, \hat{\eta}) (\hat{v}_j - v_j^*) = \mathcal{O}_{\mathbb{P}}((s_\xi \vee s_{w_j}) \log d/n). \quad (51)$$

Proof Let $\hat{\Delta}_{\mathbf{w}_j} := \hat{\mathbf{w}}_j - \mathbf{w}_j^*$ and $\hat{\Delta}_{\mathbf{v}_j} := \hat{\mathbf{v}}_j - \mathbf{v}_j^*$. Since $\|\mathbf{w}_j^*\|_1 \geq \|\hat{\mathbf{w}}_j\|_1$, we have

$$\sum_{j \in \mathcal{S}_{\mathbf{w}_j}} |w_j^*| \geq \sum_{j \in \mathcal{S}_{\mathbf{w}_j}} |\hat{w}_j| + \sum_{j \in \mathcal{S}_{\hat{\mathbf{w}}_j}} |\hat{w}_j| \geq \sum_{j \in \mathcal{S}_{\mathbf{w}_j}} |w_j^*| - \sum_{j \in \mathcal{S}_{\mathbf{w}_j}} |\hat{\Delta}_{\mathbf{w}_j, j}| + \sum_{j \in \mathcal{S}_{\hat{\mathbf{w}}_j}} |\hat{w}_j|.$$

Hence

$$\sum_{j \in \mathcal{S}_{\mathbf{w}_j}} |\hat{\Delta}_{\mathbf{w}_j, j}| \geq \sum_{j \in \mathcal{S}_{\hat{\mathbf{w}}_j}} |\hat{w}_j| = \sum_{j \in \mathcal{S}_{\hat{\mathbf{w}}_j}} |\hat{\Delta}_{\mathbf{w}_j, j}|,$$

that is, $\|\hat{\Delta}_{\mathbf{w}_j, \mathcal{S}_{\hat{\mathbf{w}}_j}}\|_1 \leq \|\hat{\Delta}_{\mathbf{w}_j, \mathcal{S}_{\mathbf{w}_j}}\|_1$. Therefore, $\|\hat{\Delta}_{\mathbf{w}_j}\|_1 \leq 2\|\hat{\Delta}_{\mathbf{w}_j, \mathcal{S}_{\mathbf{w}_j}}\|_1$. Consider the following function

$$\begin{aligned} & \hat{\Delta}_{\mathbf{w}_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \hat{\Delta}_{\mathbf{w}_j} \\ &= \left[\nabla_{\theta_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \mathbf{w}_j^{*T} \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \right] \hat{\Delta}_{\mathbf{w}_j} - \left[\nabla_{\theta_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \hat{\mathbf{w}}_j^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \right] \hat{\Delta}_{\mathbf{w}_j} \\ &=: I_1 + I_2. \end{aligned}$$

According to the definition of $\hat{\Delta}_{\mathbf{w}_j}$,

$$I_2 \leq \|\nabla_{\theta_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \hat{\mathbf{w}}_j^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}})\|_\infty \|\hat{\Delta}_{\mathbf{w}_j}\|_1 \leq \lambda_{\mathbf{w}_j} \|\hat{\Delta}_{\mathbf{w}_j}\|_1.$$

Note that for I_1 ,

$$I_1 = \mathbf{v}_j^{*T} \nabla_{\xi \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \hat{\Delta}_{\mathbf{w}_j} + \mathbf{v}_j^{*T} [\nabla_{\xi \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \nabla_{\xi \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})] \hat{\Delta}_{\mathbf{w}_j} =: I_{11} + I_{12}.$$

Since

$$\mathbb{E}[\nabla_{\nu_j \xi}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \mathbf{v}_j^*] = \mathbf{I}_{\nu_j \theta_j}^* - \mathbf{I}_{\nu_j \nu_j}^* \mathbf{w}_j^* = 0,$$

we can use the similar proof of Lemma 6 to show that

$$|I_{11}| \leq \|\mathbf{v}_j^{*T} \nabla_{\xi \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\|_\infty \|\hat{\Delta}_{\mathbf{w}_j}\|_1 \lesssim \sqrt{\log d/n} \|\hat{\Delta}_{\mathbf{w}_j}\|_1.$$

For I_{22} , use similar trick as in (47) and we have

$$|I_{12}| \lesssim \hat{\Delta}_{\boldsymbol{\xi}}^T \nabla_{\xi \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \hat{\Delta}_{\mathbf{w}_j} = \hat{\Delta}_{\boldsymbol{\xi}}^T \nabla_{\xi \xi}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \hat{\Delta}_{\mathbf{v}_j}$$

By Assumption 6, $\ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})$ is positive semidefinite when n is large enough. By the Cauchy-Schwartz inequality,

$$\begin{aligned} |I_{12}| &\leq \left[\hat{\Delta}_{\boldsymbol{\xi}}^T \nabla_{\xi \xi}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \hat{\Delta}_{\boldsymbol{\xi}} \right]^{1/2} \left[\hat{\Delta}_{\mathbf{v}_j}^T \nabla_{\xi \xi}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \hat{\Delta}_{\mathbf{v}_j} \right]^{1/2} \\ &\lesssim \sqrt{s_{\boldsymbol{\xi}} \log d/n} \left[\hat{\Delta}_{\mathbf{w}_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \hat{\Delta}_{\mathbf{w}_j} \right]^{1/2}. \end{aligned}$$

Combine the bounds for I_{11}, I_{12}, I_2 ,

$$\hat{\Delta}_{\mathbf{w}_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \hat{\Delta}_{\mathbf{w}_j} \lesssim \sqrt{\log d/n} \|\hat{\Delta}_{\mathbf{w}_j}\|_1 + \sqrt{s_{\boldsymbol{\xi}} \log d/n} \left[\hat{\Delta}_{\mathbf{w}_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \hat{\Delta}_{\mathbf{w}_j} \right]^{1/2}.$$

Further notice that

$$\left| \hat{\Delta}_{w_j}^T \left[\nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) - \nabla_{\nu_j \nu_j}^2 \ell(\theta^*, \hat{\eta}) \right] \hat{\Delta}_{w_j} \right| \lesssim s_{\xi} \sqrt{\log d/n} \left| \hat{\Delta}_{w_j}^T \left[\nabla_{\nu_j \nu_j}^2 \ell(\theta^*, \hat{\eta}) \right] \hat{\Delta}_{w_j} \right|$$

using similar trick as in (47) and $\|\hat{\Delta}_{\xi}\|_1 = \mathcal{O}_{\mathbb{P}}(s_{\xi} \sqrt{\log d/n})$. Therefore,

$$\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\theta^*, \hat{\eta}) \hat{\Delta}_{w_j} (1 - C s_{\xi} \sqrt{\log d/n}) \lesssim \hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \quad (52)$$

for some constant $C > 0$. Consequently,

$$\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \lesssim \sqrt{\log d/n} \|\hat{\Delta}_{w_j}\|_1 + \sqrt{s_{\xi} \log d/n} \left[\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \right]^{1/2}. \quad (53)$$

Consider two cases. If $\left[\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \right]^{1/2} \lesssim \sqrt{s_{\xi} \log d/n}$, according to Lemma 10,

$$s_{w_j}^{-1/2} \|\hat{\Delta}_{w_j, s_{w_j}}\|_1 \lesssim \left[\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \right]^{1/2}.$$

Combine the upper and lower bounds of $\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j}$, we have $\|\hat{\Delta}_{w_j, s_{w_j}}\|_1 \lesssim \sqrt{s_{w_j} s_{\xi} \log d/n}$. Therefore, we have

$$\|\hat{\Delta}_{w_j}\|_1 \leq 2 \|\hat{\Delta}_{w_j, s_{w_j}}\|_1 \lesssim (s_{w_j} \vee s_{\xi}) \sqrt{\log d/n}.$$

If $\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \gtrsim \sqrt{s_{\xi} \log d/n}$, (53) is equivalent to

$$\left[\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \right]^{1/2} \left\{ \left[\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \right]^{1/2} - \sqrt{s_{\xi} \log d/n} \right\} \lesssim \sqrt{\log d/n} \|\hat{\Delta}_{w_j}\|_1.$$

Since $\|\hat{\Delta}_{w_j}\|_1^2 \leq 4 \|\hat{\Delta}_{w_j, s_{w_j}}\|_1^2 \leq 4 s_{w_j} \|\hat{\Delta}_{w_j, s_{w_j}}\|_2^2 \lesssim s_{w_j} \hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j}$,

$$\left[\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \right]^{1/2} - \sqrt{s_{\xi} \log d/n} \lesssim \sqrt{s_{w_j} \log d/n}.$$

Hence, we have

$$\left[\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \right]^{1/2} \lesssim \sqrt{(s_{\xi} + s_{w_j}) \log d/n}.$$

Using the same arguments as in the first case,

$$\|\hat{\Delta}_{w_j}\|_1 \lesssim (s_{w_j} + s_{\xi}) \sqrt{\log d/n}.$$

In both cases, we can conclude that $\|\hat{\Delta}_{w_j}\|_1 \lesssim (s_{w_j} \vee s_{\xi}) \sqrt{\log d/n}$ and that

$$\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{w_j} \lesssim (s_{\xi} \vee s_{w_j}) \log d/n.$$

Finally, according to (52), we have

$$\hat{\Delta}_{w_j}^T \nabla_{\nu_j \nu_j}^2 \ell(\theta^*, \hat{\eta}) \hat{\Delta}_{w_j} \lesssim (s_{\xi} \vee s_{w_j}) \log d/n.$$

■

Lemma 10 *Denote*

$$\kappa_D(s_{w_j}) = \min \left\{ \frac{s_{w_j}^{1/2} [\mathbf{w}_j^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \mathbf{w}_j]^{1/2}}{\|\mathbf{w}_j, s_{w_j}\|_1} : \mathbf{w}_j \in \mathbb{R}^{d-1} \setminus \{0\}, \|\mathbf{w}_j, s_{w_j}^c\|_1 \leq \xi \|\mathbf{w}_j, s_{w_j}\|_1 \right\},$$

where ξ is a positive constant. Under Assumptions 1-6, $\kappa_D(s_{w_j}) \geq \kappa/\sqrt{2}$ with probability tending to one.

Proof Since $\|\mathbf{v}_j, s_{w_j}\|_1 \leq s_{w_j}^{1/2} \|\mathbf{v}_j, s_{w_j}\|_2 \leq s_{w_j}^{1/2} \|\mathbf{v}_j\|_2$, we have

$$\kappa_D^2(s_{w_j}) \geq \min \left\{ \frac{\mathbf{w}_j^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2} : \mathbf{w}_j \in \mathbb{R}^{d-1} \setminus \{0\}, \|\mathbf{w}_j, s_{w_j}^c\|_1 \leq \xi \|\mathbf{w}_j, s_{w_j}\|_1 \right\}.$$

Note that

$$\frac{\mathbf{w}_j^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2} = \frac{\mathbf{w}_j^T \nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2} + \frac{\mathbf{w}_j^T [\nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})] \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2}$$

and

$$\mathbf{w}_j^T [\nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})] \mathbf{w}_j = \mathcal{O}_{\mathbb{P}}(s_{\xi} \sqrt{\log d/n}) = o_{\mathbb{P}}(1)$$

by similar proof as (48). Therefore, with probability tending to one,

$$\begin{aligned} \frac{\mathbf{w}_j^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2} &\geq \frac{3}{4} \frac{\mathbf{w}_j^T \nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2} \\ &= \frac{3}{4} \frac{\mathbf{w}_j^T \mathbf{I}_{\nu_j \nu_j} \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2} + \frac{3}{4} \frac{\mathbf{w}_j^T [\nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbf{I}_{\nu_j \nu_j}] \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2} \\ &\geq \frac{3}{4} \lambda_{\min}(\mathbf{I}_{\nu_j \nu_j}) - \frac{3}{4} \left| \frac{\mathbf{w}_j^T [\nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbf{I}_{\nu_j \nu_j}] \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2} \right| \\ &\geq \frac{3}{4} \left[\kappa^2 - \frac{\|\mathbf{w}_j\|_1^2 \|\nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbf{I}_{\nu_j \nu_j}\|_{\infty}}{\|\mathbf{w}_j\|_2^2} \right] \end{aligned}$$

In addition, from

$$\|\mathbf{w}_j\|_1^2 \leq (\xi + 1)^2 \|\mathbf{w}_j, s_{w_j}\|_1^2 \leq s_{w_j} (\xi + 1)^2 \|\mathbf{w}_j\|_2^2$$

we get

$$\frac{\mathbf{w}_j^T \nabla_{\nu_j \nu_j}^2 \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) \mathbf{w}_j}{\|\mathbf{w}_j\|_2^2} \geq \frac{3}{4} \left[\kappa^2 - s_{w_j} (\xi + 1)^2 \|\nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbf{I}_{\nu_j \nu_j}\|_{\infty} \right].$$

Similar to the proof of Lemma 6, we can obtain

$$\|\nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbf{I}_{\nu_j \nu_j}\|_{\infty} = \mathcal{O}_{\mathbb{P}}(\sqrt{\log d/n}).$$

Hence we have $\|\nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbf{I}_{\nu_j \nu_j}\|_\infty = o_{\mathbb{P}}(1)$ by Assumption 5. When n is large enough,

$$\|\nabla_{\nu_j \nu_j}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \mathbf{I}_{\nu_j \nu_j}\|_\infty \leq \kappa^2 / [3(\xi + 1)^2].$$

Therefore, $\kappa_D(s_{w_j}) \geq \kappa / \sqrt{2}$ with probability tending to one. \blacksquare

Lemma 11 (Local smoothness conditions on the loss function) *Based on Assumptions 1-3 and 5, we have*

$$\mathbf{v}_j^{*T} [\nabla_{\boldsymbol{\xi}} \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)] = \mathcal{O}_{\mathbb{P}}((s_{\boldsymbol{\xi}} \vee s_{w_j}) \log d/n), \quad (54)$$

$$(\hat{\mathbf{v}}_j - \mathbf{v}_j^*)^T [\nabla_{\boldsymbol{\xi}} \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})] = \mathcal{O}_{\mathbb{P}}((s_{\boldsymbol{\xi}} \vee s_{w_j}) \log d/n). \quad (55)$$

The same results hold for $\hat{\boldsymbol{\theta}}_0 = (\theta_0, 0, \hat{\boldsymbol{\nu}}_j^T)^T$, where θ_j is the parameter we are interested in.

Proof Using the similar proof as (47),

$$\begin{aligned} & \left| \mathbf{v}_j^{*T} [\nabla_{\boldsymbol{\xi}} \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)] \right| \\ &= \left| \mathbf{v}_j^{*T} [\nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\tilde{q}\hat{\boldsymbol{\theta}} + (1 - \tilde{q})\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\hat{\Delta}_{\boldsymbol{\xi}} - \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\hat{\Delta}_{\boldsymbol{\xi}}] \right| \\ &\lesssim \|\mathbf{v}_j^*\|_1 \|\hat{\Delta}_{\boldsymbol{\xi}}^T \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\hat{\Delta}_{\boldsymbol{\xi}}\|_\infty \lesssim (s_{\boldsymbol{\xi}} \vee s_{w_j}) \log d/n \end{aligned}$$

for some $\tilde{q} \in [0, 1]$ by (43). For (55), we have

$$\begin{aligned} & \left| (\hat{\mathbf{v}}_j - \mathbf{v}_j^*)^T [\nabla_{\boldsymbol{\xi}} \ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) - \nabla_{\boldsymbol{\xi}} \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})] \right| \\ &= \left| \hat{\Delta}_{v_j} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\tilde{q}\hat{\boldsymbol{\theta}} + (1 - \tilde{q})\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\hat{\Delta}_{\boldsymbol{\xi}} \right| \\ &\leq \left| \hat{\Delta}_{v_j} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\hat{\Delta}_{\boldsymbol{\xi}} \right| + \left| \hat{\Delta}_{v_j} [\nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\tilde{q}\hat{\boldsymbol{\theta}} + (1 - \tilde{q})\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})]\hat{\Delta}_{\boldsymbol{\xi}} \right| \end{aligned}$$

by Taylor's expansion. For the first term on the right-hand side of the inequality, the Cauchy-Schwartz inequality yields

$$\begin{aligned} & \left| \hat{\Delta}_{v_j} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\hat{\Delta}_{\boldsymbol{\xi}} \right| \\ &\leq \left| \hat{\Delta}_{v_j} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\hat{\Delta}_{v_j} \right|^{1/2} \left| \hat{\Delta}_{\boldsymbol{\xi}} \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\hat{\Delta}_{\boldsymbol{\xi}} \right|^{1/2} \\ &\lesssim \sqrt{s_{\boldsymbol{\xi}} \log d/n} \sqrt{(s_{\boldsymbol{\xi}} \vee s_{w_j}) \log d/n} = \mathcal{O}_{\mathbb{P}}((s_{\boldsymbol{\xi}} \vee s_{w_j}) \log d/n) \end{aligned}$$

by (43) and (50) and the positive semidefiniteness of $\ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})$. For the second term on the right-hand side, similar proof as (47) yields

$$\begin{aligned} & \left| \hat{\Delta}_{v_j} [\nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\tilde{q}\hat{\boldsymbol{\theta}} + (1 - \tilde{q})\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}}) - \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})]\hat{\Delta}_{\boldsymbol{\xi}} \right| \\ &\lesssim \|\hat{\Delta}_{v_j}\|_1 \|\hat{\Delta}_{\boldsymbol{\xi}}^T \nabla_{\boldsymbol{\xi}\boldsymbol{\xi}}^2 \ell(\boldsymbol{\theta}^*, \hat{\boldsymbol{\eta}})\hat{\Delta}_{\boldsymbol{\xi}}\|_\infty \\ &\lesssim (s_{\boldsymbol{\xi}} \vee s_{w_j}) \sqrt{\log d/n} ((s_{\boldsymbol{\xi}} \vee s_{w_j}) \log d/n) \end{aligned}$$

by (43) and (49). ■

Finally, we give the proof of Theorem 2.

Proof For the assumptions needed in Ning and Liu (2017, Theorem 3.2), Lemmas 1 and 9 satisfy Assumption 1, Lemma 6 satisfies Assumption 2, Lemma 11 satisfies Assumption 3, and Lemma 7 satisfies Assumption 4. Now we only need to verify that $\hat{I}_{\theta_j|\nu_j} - I_{\theta_j|\nu_j}^* = o_{\mathbb{P}}(1)$.

First note that

$$\hat{I}_{\theta_j|\nu_j} = \nabla_{\theta_j\theta_j}^2 \ell(\hat{\theta}, \hat{\eta}) - \hat{w}_j^T \nabla_{\nu_j\theta_j}^2 \ell(\hat{\theta}, \hat{\eta})$$

and

$$I_{\theta_j|\nu_j}^* = I_{\theta_j\theta_j}^* - w_j^{*T} \mathbf{I}_{\nu_j\theta_j}^*, \quad \text{where } w_j^* = \mathbf{I}_{\nu_j\nu_j}^{*-1} \mathbf{I}_{\nu_j\theta_j}^*.$$

For the second part of $\hat{I}_{\theta_j|\nu_j}$ and $I_{\theta_j|\nu_j}^*$,

$$\begin{aligned} & \hat{w}_j^T \nabla_{\nu_j\theta_j}^2 \ell(\hat{\theta}, \hat{\eta}) - w_j^{*T} \mathbf{I}_{\nu_j\theta_j}^* \\ & \leq \left| (\hat{w}_j - w_j^*)^T \nabla_{\nu_j\theta_j}^2 \ell(\hat{\theta}, \hat{\eta}) \right| + \left| w_j^{*T} [\nabla_{\nu_j\theta_j}^2 \ell(\hat{\theta}, \hat{\eta}) - \nabla_{\nu_j\theta_j}^2 \ell(\theta^*, \hat{\eta})] \right| \\ & \quad + \left| w_j^{*T} [\nabla_{\nu_j\theta_j}^2 \ell(\theta^*, \hat{\eta}) - \mathbf{I}_{\nu_j\theta_j}^*] \right| \\ & =: J_1 + J_2 + J_3. \end{aligned}$$

With $e = [1, 0, \dots, 0]$,

$$J_1 = \left| \hat{\Delta}_{v_j}^T \nabla_{\xi\xi}^2 \ell(\hat{\theta}, \hat{\eta}) e \right| \leq \left\| \hat{\Delta}_{v_j}^T \nabla_{\xi\xi}^2 \ell(\hat{\theta}, \hat{\eta}) \hat{\Delta}_{v_j} \right\|_2 \left\| e^T \nabla_{\xi\xi}^2 \ell(\hat{\theta}, \hat{\eta}) e \right\| \lesssim \sqrt{(s_{\xi} \vee s_{w_j}) \log d/n}$$

by the Cauchy-Schwartz inequality, the bound (51) and the boundedness of the variables. Using similar arguments as in (47),

$$J_2 \lesssim \left| v_j^{*T} \nabla_{\xi\xi}^2 \ell(\theta^*, \hat{\eta}) (\hat{\theta} - \theta^*) \right|.$$

Similarly, with the Cauchy-Schwartz inequality, the bound (43) and the boundedness of the variables,

$$J_2 \leq \left\| (\hat{\theta} - \theta^*)^T \nabla_{\xi\xi}^2 \ell(\hat{\theta}, \hat{\eta}) (\hat{\theta} - \theta^*) \right\|_2 \lesssim \sqrt{s_{\xi} \log d/n}$$

Using similar arguments as the proof of (40), we have that

$$J_3 \lesssim \sqrt{\log d/n}.$$

In addition, similar arguments also imply that

$$|\nabla_{\theta_j\theta_j}^2 \ell(\hat{\theta}, \hat{\eta}) - I_{\theta_j\theta_j}^*| \lesssim \sqrt{\log d/n}.$$

Therefore, $|\hat{I}_{\theta_j|\nu_j} - I_{\theta_j|\nu_j}^*| \lesssim \sqrt{(s_{\xi} \vee s_{w_j}) \log d/n} = o_{\mathbb{P}}(1)$ and thus $\sqrt{n}(\tilde{\theta}_j - \theta_j^*) I_{\theta_j|\nu_j}^* \sqrt{\sigma_j^*} \Rightarrow N(0, 1)$. ■

References

- András Antos, Csaba Szepesvári, and Rémi Munos. Fitted Q-iteration in continuous action-space mdps. In *Advances in Neural Information Processing Systems*, volume 20, 2007.
- Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.
- Audrey Boruvka, Daniel Almirall, Katie Witkiewitz, and Susan A Murphy. Assessing time-varying causal effect moderation in mobile health. *Journal of the American Statistical Association*, 113(523):1112–1121, 2018.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018a.
- Victor Chernozhukov, Denis Nekipelov, Vira Semenova, and Vasilis Syrgkanis. Plug-in regularized estimation of high dimensional parameters in nonlinear semiparametric models. Technical report, Centre for Microdata Methods and Practice, Institute for Fiscal Studies, 2018b.
- Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- Thomas S Ferguson. *A Course in Large Sample Theory*. Routledge, 2017.
- Daiqi Gao, Yufeng Liu, and Donglin Zeng. Non-asymptotic properties of individualized treatment rules from sequentially rule-adaptive trials. *Journal of Machine Learning Research*, 23(250):1–42, 2022.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2009.
- Yuchen Hu and Stefan Wager. Off-policy evaluation in partially observed Markov decision processes under sequential ignorability. *Annals of Statistics*, 51(4):1561–1585, 2023.
- X. Jessie Jeng, Wenbin Lu, and Huimin Peng. High-dimensional inference for personalized treatment decision. *Electronic Journal of Statistics*, 12(1):2074, 2018.
- Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661. PMLR, 2016.
- Nathan Kallus and Masatoshi Uehara. Double reinforcement learning for efficient off-policy evaluation in Markov decision processes. *Journal of Machine Learning Research*, 21(167), 2020.
- Nathan Kallus and Masatoshi Uehara. Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning. *Operations Research*, 70(6):3282–3302, 2022.

- Michael R Kosorok and Eric B Laber. Precision medicine. *Annual Review of Statistics and Its Application*, 6(1):263–286, 2019.
- Peter D Lax and Maria Shea Terrell. *Calculus with Applications*, volume 4. Springer, 2014.
- Muxuan Liang, Young-Geun Choi, Yang Ning, Maureen A Smith, and Ying-Qi Zhao. Estimation and inference on high-dimensional individualized treatment rule in observational data using split-and-pooled de-correlated score. *Journal of Machine Learning Research*, 23(262):1–65, 2022.
- Peng Liao, Predrag Klasnja, and Susan Murphy. Off-policy estimation of long-term average outcomes with applications to mobile health. *Journal of the American Statistical Association*, 116(533):382–391, 2021.
- Peng Liao, Zhengling Qi, Runzhe Wan, Predrag Klasnja, and Susan A Murphy. Batch policy learning in average reward Markov decision processes. *Annals of Statistics*, 50(6):3364–3387, 2022.
- Qiang Liu and Jason Lee. Black-box importance sampling. In *Artificial Intelligence and Statistics*, pages 952–961. PMLR, 2017.
- Qiang Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems*, pages 5356–5366, 2018a.
- Ying Liu, Yuanjia Wang, Michael R Kosorok, Yingqi Zhao, and Donglin Zeng. Augmented outcome-weighted learning for estimating optimal dynamic treatment regimens. *Statistics in Medicine*, 37(26):3776–3788, 2018b.
- Daniel J Lockett, Eric B Laber, Anna R Kahkoska, David M Maahs, Elizabeth Mayer-Davis, and Michael R Kosorok. Estimating dynamic treatment regimes in mobile health using V-learning. *Journal of the American Statistical Association*, 115(530):692–706, 2020.
- Alexander R Luedtke and Mark J Van Der Laan. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *Annals of Statistics*, 44(2):713, 2016.
- Haixu Ma, Donglin Zeng, and Yufeng Liu. Learning individualized treatment rules with many treatments: A supervised clustering approach using adaptive fusion. In *Advances in Neural Information Processing Systems*, pages 15956–15969, 2022.
- Haixu Ma, Donglin Zeng, and Yufeng Liu. Learning optimal group-structured individualized treatment rules with many treatments. *Journal of Machine Learning Research*, 24(102):1–48, 2023.
- Cindy Marling and Razvan Bunescu. The OhioT1DM dataset for blood glucose level prediction: Update 2020. In *CEUR Workshop Proceedings*, volume 2675, page 71. NIH Public Access, 2020.

- Erica EM Moodie, Bibhas Chakraborty, and Michael S Kramer. Q-learning for estimating optimal dynamic treatment rules from observational data. *Canadian Journal of Statistics*, 40(4):629–645, 2012.
- Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003.
- Susan A Murphy. A generalization error for Q-learning. *Journal of Machine Learning Research*, 6, 2005.
- Xinkun Nie, Emma Brunskill, and Stefan Wager. Learning when-to-treat policies. *Journal of the American Statistical Association*, 116(533):392–409, 2021.
- Yang Ning and Han Liu. A general theory of hypothesis tests and confidence regions for sparse high dimensional models. *Annals of Statistics*, 45(1):158–195, 2017.
- Doina Precup. *Temporal abstraction in reinforcement learning*. PhD thesis, University of Massachusetts Amherst, 2000.
- Zhengling Qi, Dacheng Liu, Haoda Fu, and Yufeng Liu. Multi-armed angle-based direct learning for estimating optimal individualized treatment rules with various outcomes. *Journal of the American Statistical Association*, 115(530):678–691, 2020.
- James M Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*, volume 1999, pages 6–10. Indianapolis, IN, 2000.
- Chengchun Shi, Alin Fan, Rui Song, and Wenbin Lu. High-dimensional A-learning for optimal dynamic treatment regimes. *Annals of Statistics*, 46(3):925, 2018.
- Chengchun Shi, Wenbin Lu, and Rui Song. Breaking the curse of nonregularity with sub-agging: inference of the mean outcome under optimal treatment regimes. *Journal of Machine Learning Research*, 21, 2020a.
- Chengchun Shi, Runzhe Wan, Rui Song, Wenbin Lu, and Ling Leng. Does the Markov decision process fit the data: Testing for the Markov property in sequential decision making. In *International Conference on Machine Learning*, pages 8807–8817. PMLR, 2020b.
- Chengchun Shi, Runzhe Wan, Victor Chernozhukov, and Rui Song. Deeply-debiased off-policy interval estimation. In *International Conference on Machine Learning*, pages 9580–9591, 2021.
- Chengchun Shi, Sheng Zhang, Wenbin Lu, Rui Song, et al. Statistical inference of the value function for reinforcement learning in infinite-horizon settings. *Journal of the Royal Statistical Society Series B*, 84(3):765–793, 2022.
- Suvrit Sra, Sebastian Nowozin, and Stephen J Wright. *Optimization for Machine Learning*. MIT Press, 2011.

- Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.
- Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148. PMLR, 2016.
- Philip S Thomas. *Safe Reinforcement Learning*. PhD thesis, University of Massachusetts Amherst, 2015.
- Aad van der Vaart and Jon A Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Science & Business Media, 1996.
- Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, 100(3):681–694, 2013.
- Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *Annals of Statistics*, 38(2):894–942, 2010.
- Cun-Hui Zhang and Stephanie S Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):217–242, 2014.
- Ying-Qi Zhao, Donglin Zeng, Eric B Laber, and Michael R Kosorok. New statistical learning methods for estimating optimal dynamic treatment regimes. *Journal of the American Statistical Association*, 110(510):583–598, 2015.
- Yufan Zhao, Donglin Zeng, Mark A Socinski, and Michael R Kosorok. Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. *Biometrics*, 67(4):1422–1433, 2011.
- Wensheng Zhu, Donglin Zeng, and Rui Song. Proper inference for value function in high-dimensional Q-learning for dynamic treatment regimes. *Journal of the American Statistical Association*, 114(527):1404–1417, 2019.