

Policy Learning for Optimal Dynamic Treatment Regimes with Observational Data

Shosei Sakaguchi*

December 10, 2024

Abstract

Public policies and medical interventions often involve dynamics in their treatment assignments, where individuals receive a series of interventions over multiple stages. We study the statistical learning of optimal dynamic treatment regimes (DTRs) that guide the optimal treatment assignment for each individual at each stage based on the individual's evolving history. We propose a doubly robust, classification-based approach to learning the optimal DTR using observational data under the assumption of sequential ignorability. This approach learns the optimal DTR through backward induction. At each step, it constructs an augmented inverse probability weighting (AIPW) estimator of the policy value function and maximizes it to learn the optimal policy for the corresponding stage. We show that the resulting DTR can achieve an optimal convergence rate of $n^{-1/2}$ for welfare regret under mild convergence conditions on estimators of the nuisance components.

Keywords: Dynamic treatment regime; Policy learning; Double robustness; Double machine learning; Backward induction.

*Faculty of Economics, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan.
Email: sakaguchi@e.u-tokyo.ac.jp.

1 Introduction

Public policies and medical interventions often involve dynamics in their treatment assignments. For example, some job training programs offer participants training sessions over multiple stages (e.g., [Lechner, 2009](#); [Rodríguez et al., 2022](#)). In clinical medicine, physicians sequentially administer treatments, adapting to patients’ evolving medical conditions (e.g., [Wang et al., 2012](#); [Pelham Jr et al., 2016](#)). Multi-stage treatment assignments are also common in educational programs spanning multiple grades (e.g., [Krueger, 1999](#); [Ding and Lehrer, 2010](#)) and dynamic marketing strategies (e.g., [Liu, 2023](#)), among other examples.

This study focuses on the optimal allocation of sequential treatments ([Robins, 1986](#)), where individuals receive interventions over multiple stages. In this context, treatment effects at each stage are often heterogeneous, depending on prior treatments and related states. Therefore, adapting treatment allocation based on evolving information can substantially enhance the welfare gains of multi-stage interventions.

We study statistical learning for optimal sequential treatment assignment using data from quasi-experimental or observational studies. Throughout this paper, we assume sequential ignorability ([Robins, 1997](#)), meaning that treatment assignment at each stage is independent of potential outcomes, conditional on the history of prior treatments and observed states. Under this assumption, we develop a method to learn the optimal Dynamic Treatment Regime (DTR), a sequence of stage-specific policies that determines the optimal treatment for each individual at each stage, based on their history up to that point.

In developing our approach, we build on recent advances in doubly robust policy learning ([Athey and Wager, 2021](#); [Zhou et al., 2023b](#)) and extend them to dynamic settings. We propose a doubly robust, classification-based method to learn the optimal DTR using backward induction, which sequentially estimates the optimal policy from the final to the first stage. At each step of the backward induction, the method constructs an augmented inverse probability weighting (AIPW) estimator for the policy value function by combining estimators of propensity scores and action value functions (Q-functions) while using cross-fitting.¹ The Q-functions can be estimated via fitted Q-evaluation ([Munos and Szepesvári, 2008](#); [Fonteneau et al., 2013](#); [Le et al., 2019](#)), a method for offline policy eval-

¹The Q-function for each stage of a fixed DTR is the conditional mean outcome given the treatment and history, assuming future treatments follow the fixed DTR.

uation in reinforcement learning. At each step, the optimal policy for the corresponding stage is estimated by maximizing the estimated policy value function over a pre-specified class of stage-specific policies. This procedure produces the estimated DTR as a sequence of the estimated policies across all stages.

The proposed approach is computationally feasible due to step-wise backward optimization, which is a considerable advantage given the computational challenge of optimizing DTRs over multiple stages. Additionally, it does not rely on structural assumptions, such as Markov decision processes. Leveraging a classification-based framework (Kitagawa and Tetenov, 2018), the approach can enhance the interpretability of DTRs by incorporating interpretable policy classes, such as decision trees, for each stage. Moreover, the approach can encompass various dynamic treatment problems, including optimal stopping/starting problems, by appropriately constraining the class of feasible DTRs.²

We study the statistical properties of the proposed approach in terms of welfare regret, defined as the average outcome loss of the estimated DTR relative to the optimal one. This is, however, a nontrivial task because stage-specific policies within the DTR are estimated sequentially, rather than simultaneously, and state variables are influenced by past treatments. The main contribution of this paper is to establish the convergence rate for welfare regret, linking it to the convergence rates of the nuisance component estimators (propensity scores and Q-functions) in terms of the mean-squared-error (MSE). Our key result identifies conditions on the nuisance component estimators and the class of DTRs under which the resulting DTR attains the minimax optimal convergence rate of $n^{-1/2}$ for regret. For example, if all nuisance components are estimated with an MSE convergence rate of $n^{-1/4}$, which is attainable by many machine learning methods under structured assumptions, the resulting DTR can achieve regret convergence to zero at the minimax optimal rate of $n^{-1/2}$. This result is comparable to those of Athey and Wager (2021) and Zhou et al. (2023b), who study doubly robust policy learning in single-stage settings, and aligns with the spirit of double machine learning (Chernozhukov et al., 2018).

We illustrate the proposed method with an empirical application to data from Project STAR (e.g., Krueger, 1999), where we learn the optimal DTR for assigning each student to either a regular-size class with a teacher aide or a small-size class without one during their early education (kindergarten and grade 1). The estimated DTR uses information on students' intermediate academic performance to determine the optimal class type for

²An important example of an optimal stopping problem in economics is unemployment insurance programs that reduce benefit levels after a certain duration (e.g., Meyer, 1995; Kolsrud et al., 2018)

a subsequent grade. Our empirical results show that the optimal DTR leads to better academic outcomes for students compared with uniform class-type allocations.

Related Literature

Although numerous studies have investigated the statistical decision/learning of treatment choice, most have focused on the single-stage problem.³ Among these works, this study is most relevant to [Athey and Wager \(2021\)](#) and [Zhou et al. \(2023b\)](#). They propose doubly robust policy learning in single-stage settings and show that the $n^{-1/2}$ -upper bound on the regret can be achieved even in the observational data setting. This study seeks to extend their approach and results to the dynamic treatment problem.⁴

There is an expanding literature on the estimation of optimal DTRs, with many methods proposed.⁵ Offline Q-learning ([Watkins and Dayan, 1992](#)) is arguably the most popular method for estimating optimal DTRs (e.g., [Murphy, 2005](#); [Moodie et al., 2012](#); [Zhang et al., 2018](#)).⁶ [Murphy \(2005\)](#) shows that the performance of the DTR obtained via Q-learning depends on the accuracy of the Q-function estimates; if the estimated Q-functions are not close to the true ones, the resulting DTR can be far from optimal. Our approach also estimates Q-functions as part of its procedure. However, leveraging propensity score models, it is more robust and accurate than Q-learning.

This study is also related to the classification-based, inverse probability weighting approach for estimating optimal DTRs (e.g., [Zhao et al., 2015](#); [Sakaguchi, 2024](#)). This approach uses inverse propensity weighted outcomes to estimate the value function of a DTR and maximizes it to estimate the optimal DTR over a pre-specified class of DTRs. However, the use of inverse probability weighted outcomes sometimes leads to excessively high variance, resulting in suboptimal DTRs ([Doroudi et al., 2018](#)). Our approach enhances the power of this method by incorporating models of Q-functions.

Doubly robust estimation for optimal DTRs has also been proposed by [Zhang et al.](#)

³A partial list includes [Manski \(2004\)](#), [Hirano and Porter \(2009\)](#), [Stoye \(2009; 2012\)](#), [Qian and Murphy \(2011\)](#), [Bhattacharya and Dupas \(2012\)](#), [Tetenov \(2012\)](#), [Zhao et al. \(2012\)](#), [Kitagawa and Tetenov \(2018\)](#), [Athey and Wager \(2021\)](#), [Mbakop and Tabor-Meehan \(2021\)](#), [Kitagawa et al. \(2023\)](#), [Zhou et al. \(2023b\)](#), and [Viviano \(2024\)](#), among others.

⁴To be precise, this study extends the results of [Zhou et al. \(2023b\)](#) rather than those of [Athey and Wager \(2021\)](#), as it involves multiple treatments at each stage and relies on a similar complexity condition on the policy class as in [Zhou et al. \(2023b\)](#).

⁵[Chakraborty and Murphy \(2014\)](#), [Laber et al. \(2014\)](#), [Kosorok and Laber \(2019\)](#), and [Li et al. \(2023\)](#) provide reviews of this literature.

⁶Q-learning is a sequential approach that estimates the optimal Q-functions and optimal policies from the final stage to the first stage through backward induction.

(2013), Wallace and Moodie (2015), and Ertefaie et al. (2021). Zhang et al. (2013) suggest estimating the optimal DTR by maximizing an AIPW estimator of the welfare function over an entire class of DTRs. However, this approach faces computational challenges for two reasons: (i) the nuisance components to be estimated depend on each specific DTR, and (ii) the method maximizes the estimated welfare function simultaneously over the entire class of DTRs. Our approach resolves these issues because (i) the nuisance components depend only on the estimated policies for future stages, and (ii) the optimal DTR is estimated through stage-wise backward optimization. Section 5 provides further details on this comparison. Wallace and Moodie (2015) develop a doubly robust estimation method based on Q-learning and G-estimation (Robins, 2004). Ertefaie et al. (2021) propose a doubly robust approach for Q-learning and investigate the statistical properties of the estimated parameters in the Q-functions. By contrast, our focus is on the statistical properties of welfare regret for the estimated DTR. In the context of policy learning for optimal stopping/starting problems, Nie et al. (2021) develop a doubly robust learning approach with computational feasibility and show upper bounds on the associated regret. Our framework encompasses this problem as a specific case.

This study also relates to the literature on (offline) reinforcement learning in terms of multi-stage decision problems.⁷ However, most works assume a Markov decision process, which this study does not rely on. In the context of non-Markov decision processes, Jiang and Li (2016), Thomas and Brunskill (2016), and Kallus and Uehara (2020) propose doubly robust methods for evaluating DTRs, but they do not focus on optimizing DTRs.

Finally, in the econometric literature on dynamic treatment analysis, Heckman and Navarro (2007) and Heckman et al. (2016) use exclusion restrictions to identify the average dynamic treatment effects, but their focus is not on the identification of optimal DTRs. Han (2023) proposes a method to characterize the sharp partial ordering of the counterfactual welfares of DTRs in an instrumental variable setting. Ida et al. (2024) empirically show that the optimal DTR outperforms optimal static targeting policies in the context of energy-saving rebate programs.

Structure of the Paper

The remainder of the paper is organized as follows. Section 2 outlines the dynamic treatment choice problem. Section 3 presents the doubly robust approach for learning the

⁷This study differs from the multi-armed bandit problem, as the latter involves online learning, whereas this study focuses on offline learning.

optimal DTRs through backward induction. Section 4 shows the statistical properties of the proposed approach. Section 5 compares the proposed approach with that of Zhang et al. (2013). Section 6 presents a simulation study to evaluate the finite sample performance of the proposed approach. Section 7 shows the empirical application results. Appendix A includes the proof of the main theorem along with some auxiliary lemmas.

2 Setup

Section 2.1 introduces the dynamic treatment framework, following the dynamic counterfactual outcome framework of Robins (1986, 1997) and Murphy (2003). Subsequently, we define the dynamic treatment choice problem in Section 2.2.

2.1 Dynamic Treatment Framework

We consider a fixed number of stages, T ($< \infty$), for multiple treatment assignments. Let $\mathcal{A}_t \equiv \{0, \dots, d_t - 1\}$ ($t = 1, \dots, T$) denote the set of possible treatment arms at stage t , where d_t (≥ 2) represents the number of treatment arms at stage t and may vary across stages. We observe the assigned treatment $A_t \in \mathcal{A}_t$ for each individual in each stage t . Let S_t be a vector of the state variables observed prior to treatment assignment in stage t , which may depend on the past treatments. In each stage t , we observe the outcome Y_t after the treatment intervention. S_t ($t \geq 2$) may contain the previous outcomes (Y_1, \dots, Y_{t-1}) . Throughout this paper, for any time-dependent object V_t , we denote by $\underline{V}_t \equiv (V_1, \dots, V_t)$ the history of the object up to stage t , and by $\underline{V}_{s:t} \equiv (V_s, \dots, V_t)$, for $s \leq t$, the partial history of the object from stage s up to stage t . For example, \underline{A}_t denotes the treatment history up to stage t .

Let $Z \equiv (\underline{A}_T, \underline{S}_T, \underline{Y}_T)$ represent the vector of all observed variables. We define the history at stage t as $H_t \equiv (\underline{A}_{t-1}, \underline{S}_t)$, which is the information available when the policymaker selects the treatment intervention at stage t . Note that $H_1 = (S_1)$, where S_1 represents individual characteristics observed prior to the beginning of the sequential treatment intervention. We denote the supports of H_t and Z by \mathcal{H}_t and \mathcal{Z} , respectively.

To formalize our results, we employ the framework of dynamic potential outcomes (Robins, 1986; Hernán et al., 2001; Murphy, 2003). Let $\underline{A}_t \equiv \mathcal{A}_1 \times \dots \times \mathcal{A}_t$. We define $Y_t(a_t)$ as the potential outcome of $a_t \in \mathcal{A}_t$ for stage t , which represents the outcome realized at stage t when the treatment history up to that stage corresponds to \underline{a}_t . We

assume that the outcome at each stage is not influenced by treatments in future stages. Since the state variables S_t may also depend on past treatments, we define the potential state variables as $S_t(\underline{a}_{t-1})$ for each $t \geq 2$ and $\underline{a}_{t-1} \in \underline{\mathcal{A}}_{t-1}$. For $t = 1$, we set $S_1(\underline{a}_0) = S_1$. The observed outcomes and state variables are thus defined as $Y_t \equiv Y_t(\underline{A}_t)$ and $S_t \equiv S_t(\underline{A}_{t-1})$, respectively.

Let $\underline{S}_t(\underline{a}_{t-1}) \equiv (S_1, S_2(\underline{a}_1), \dots, S_t(\underline{a}_{t-1}))$. We define the vector $H_t(\underline{a}_{t-1}) \equiv (\underline{a}_{t-1}, \underline{S}_t(\underline{a}_{t-1}))$ as the potential history realized when the prior treatments are \underline{a}_{t-1} . For $t = 1$, we set $H_1(\underline{a}_0) = H_1$. The observed history is then defined as $H_t \equiv H_t(\underline{A}_{t-1})$. We denote by P the distribution of all underlying variables $(\underline{A}_T, \{\underline{S}_T(\underline{a}_{T-1})\}_{\underline{a}_{T-1} \in \underline{\mathcal{A}}_{T-1}}, \{\underline{Y}_T(\underline{a}_T)\}_{\underline{a}_T \in \underline{\mathcal{A}}_T})$, where $\underline{Y}_T(\underline{a}_T) \equiv (Y_1(\underline{a}_1), Y_2(\underline{a}_2), \dots, Y_T(\underline{a}_T))$.

From an observational study, we observe $Z_i \equiv (\underline{A}_{i,T}, \underline{S}_{i,T}, \underline{Y}_{i,T})$ for individuals $i = 1, \dots, n$, where $\underline{A}_{i,T} = (A_{i,1}, \dots, A_{i,T})$, $\underline{S}_{i,T} = (S_{i,1}, \dots, S_{i,T})$, and $\underline{Y}_{i,T} = (Y_{i,1}, \dots, Y_{i,T})$. The observed outcome $Y_{i,t}$ and state variables $S_{i,t}$ are defined as $Y_{i,t} \equiv Y_{i,t}(\underline{A}_{i,t})$ and $S_{i,t} \equiv S_{i,t}(\underline{A}_{i,t-1})$, respectively, with $Y_{i,t}(\underline{a}_t)$ and $S_{i,t}(\underline{a}_{t-1})$ being the potential outcome and state variables, respectively, for individual i at stage t . Let $\underline{S}_{i,T}(\underline{a}_{T-1}) \equiv (S_{i,1}, S_{i,2}(\underline{a}_1), \dots, S_{i,T}(\underline{a}_{T-1}))$ and $\underline{Y}_{i,T}(\underline{a}_T) \equiv (Y_{i,2}(\underline{a}_1), \dots, Y_{i,T}(\underline{a}_T))$. We assume that the vectors of random variables $(\underline{A}_{i,T}, \{\underline{S}_{i,T}(\underline{a}_{T-1})\}_{\underline{a}_{T-1} \in \underline{\mathcal{A}}_{T-1}}, \{\underline{Y}_{i,T}(\underline{a}_T)\}_{\underline{a}_T \in \underline{\mathcal{A}}_T})$, for $i = 1, \dots, n$, are independent and identically distributed (i.i.d) under the distribution P . We denote by $H_{i,t} \equiv (\underline{A}_{i,t-1}, \underline{S}_{i,t})$ the history of the i -th individual at stage t .

Define $e_t(h_t, a_t) \equiv \Pr(A_t = a_t \mid H_t = h_t)$ as the propensity score of treatment a_t at stage t with the history h_t . In the observational data setting we study, the propensity scores are unknown to the analyst. This contrasts with the experimental data setting, in which the propensity scores are known from the experimental design.⁸

Throughout the paper, we suppose that the underlying distribution P satisfies the following assumptions.

Assumption 2.1 (Sequential Ignorability). *For any $t = 1, \dots, T$ and $\underline{a}_T \in \underline{\mathcal{A}}_T$,*

$$\{Y_t(\underline{a}_t), \dots, Y_T(\underline{a}_T), S_{t+1}(\underline{a}_t), \dots, S_T(\underline{a}_{T-1})\} \perp\!\!\!\perp A_t \mid H_t \quad a.s.$$

Assumption 2.2 (Bounded Outcomes). *There exists $M < \infty$ such that the support of $Y_t(\underline{a}_t)$ is contained in $[-M/2, M/2]$ for all $t \in \{1, \dots, T\}$ and $\underline{a}_t \in \underline{\mathcal{A}}_t$.*

⁸Even when the data are obtained from a sequential randomized trial, the propensity scores may be unknown due to non-compliance with assigned treatments or attrition over stages.

Assumption 2.1 is what is called a dynamic unconfoundedness assumption or sequential conditional independence assumption elsewhere, and is commonly used in the literature of dynamic treatment effect analysis (Robins, 1997; Murphy, 2003). This assumption implies that the treatment assignment at each stage is independent of both current and future potential outcomes and future state variables, conditional on the history up to that point. In observational studies, this assumption holds when a sufficient set of confounders is controlled at each stage. Assumption 2.2 is a common assumption in the literature on estimating optimal DTRs.

2.2 Dynamic Treatment Choice Problem

The aim of this study is to develop a method for learning optimal DTRs using data from an observational study. We define a policy for each stage t as $\pi_t : \mathcal{H}_t \mapsto \mathcal{A}_t$, a map from the history space for stage t to the treatment space for stage t . A policy π_t determines which treatment is assigned to each individual at stage t based on their history h_t . We define a DTR as $\pi \equiv (\pi_1, \dots, \pi_T)$, a sequence of stage-specific policies. The DTR guides the treatment choice for each individual from the first to the final stage, based on their evolving history up to each stage.

Given a fixed DTR π , we define the welfare of π as

$$W(\pi) \equiv E \left[\sum_{t=1}^T \sum_{\underline{a}_t \in \underline{\mathcal{A}}_t} \left(Y_t(\underline{a}_t) \cdot \prod_{s=1}^t 1\{\pi_s(H_s(\underline{a}_{s-1})) = a_s\} \right) \right].$$

This represents the expected total outcome realized when the sequential treatment assignments follow the DTR π .

We consider choosing a DTR from a pre-specified class of DTRs denoted by $\Pi \equiv \Pi_1 \times \dots \times \Pi_T$, where Π_t represents a class of policies for stage t (i.e., a class of measurable functions $\pi_t : \mathcal{H}_t \rightarrow \mathcal{A}_t$). For example, Laber and Zhao (2015), Tao et al. (2018), Sun and Wang (2021), Blumlein et al. (2022), and Zhou et al. (2023a) use a class of decision trees (Breiman et al., 1984) for Π_t , and Zhang et al. (2018) use a class of list-form policies for Π_t . These policy classes are employed to enhance the interpretability of the resulting DTRs.

Our framework can also accommodate cases where only a subset of the history, $\tilde{h}_t \subseteq h_t$, is used for treatment choice by constraining Π_t to be a class of functions of \tilde{h}_t . This is particularly important given the increasing dimension of the entire history h_t over time,

where only a sub-history may be informative for optimal treatment choice. Additionally, our framework can encompass the optimal stopping/starting problem by constraining the class Π of DTRs as follows.

Example 2.1 (Optimal Starting/Stopping Problem). *Suppose that the number of treatment arms d_t is time-invariant ($d_t = K$ for some K) and that arm 0 represents no treatment. Our framework can accommodate the optimal starting problem, where the decision-maker determines when to start assigning one of the arms $a_t \in \{1, \dots, K-1\}$ for each unit. This problem can be encompassed in our framework by constraining the class of DTRs Π such that for any $t \in \{2, \dots, T\}$ and $(\pi_t, h_t) \in \Pi_t \times \mathcal{H}_t$, $\pi_t(h_t) = a_{t-1}$ if $a_{t-1} \neq 0$. The optimal stopping problem can also be specified in a similar manner.*

An important example of an optimal stopping problem in economics is unemployment insurance programs that reduce benefit levels after a certain duration (e.g., Meyer, 1995; Kolsrud et al., 2018). In this context, a DTR determines the timing of benefit reductions for each unemployed individual.

Given a pre-specified class Π of DTRs, we impose an overlap condition on \underline{D}_T , associated with the structure of Π , as follows.

Assumption 2.3 (Overlap Condition). *There exists $\eta \in (0, 1)$ for which $\eta \leq e_t(h_t, a_t)$ holds for any $t \in \{1, \dots, T\}$ and any pair $(h_t, a_t) \in \mathcal{H}_t \times \mathcal{A}_t$ such that there exists $\pi_t \in \Pi_t$ that satisfies $\pi_t(h_t) = a_t$.*

When Π is structurally constrained, Assumption 2.3 is weaker than the common overlap condition that requires the overlap $e_t(h_t, a_t) \in (0, 1)$ for all $(h_t, a_t) \in \mathcal{H}_t \times \mathcal{A}_t$ and t . For example, in the optimal starting problem, Assumption 2.3 does not require $e_t(h_t, 0) > 0$ for any h_t such that a_s in h_t differs from 0 for some $s < t$.

The ultimate goal of our analysis is to choose an optimal DTR that maximizes the welfare $W(\cdot)$ over Π . We are especially interested in learning the optimal DTR from observational data that satisfies the sequential ignorability assumption (Assumption 2.1). The following section presents a step-wise approach to learning the optimal DTR.

3 Learning of the Optimal DTR

In this section, we propose a doubly robust approach to learning the optimal DTRs through backward induction. Section 3.1 first introduces Q-function (action-value func-

tion) and the fitted Q-evaluation, a method to estimate the Q-functions. This section also discusses identifiability of the optimal DTR through the backward-induction procedure. Section 3.2 then presents our proposed approach to learning the optimal DTRs.

3.1 Fitted Q-evaluation and Backward Induction

For any DTR π and the class Π of DTRs, we denote their partial sequences by $\pi_{s:t} \equiv (\pi_s, \dots, \pi_t)$ and $\Pi_{s:t} \equiv \Pi_s \times \dots \times \Pi_t$, respectively, for $s \leq t$.⁹ We define the policy value of $\pi_{t:T}$ for stage t as

$$V_t(\pi_{t:T}) \equiv E \left[\sum_{s=t}^T \sum_{\underline{a}_s \in \underline{\mathcal{A}}_s} \left(Y_s(\underline{a}_s) \cdot 1\{\underline{A}_{t-1} = \underline{a}_{t-1}\} \cdot \prod_{\ell=t}^s 1\{\pi_\ell(H_\ell(\underline{a}_{\ell-1})) = a_\ell\} \right) \right],$$

where we assume $1\{\underline{A}_0 = \underline{a}_0\} = 1$ for $t = 1$. $V_t(\pi_{t:T})$ represents the average total outcome from stage t to stage T that is realized when the treatment assignments before stage t follow \underline{A}_{t-1} (i.e., assignments in the observational data), and those from stage t follow $\pi_{t:T}$. With some abuse of terminology, we refer to $V_t(\pi_{t:T})$ as the policy value function of $\pi_{t:T}$. Note that $V_1(\pi_{1:T}) = W(\pi)$.

Given a fixed DTR π , we define Q-functions (state-action-value functions), recursively, as follows:

$$Q_T(h_T, a_T) \equiv E[Y_T | H_T = h_T, A_T = a_T], \quad (1)$$

$$Q_{T-1}^{\pi_T}(h_{T-1}, a_{T-1}) \equiv E[Y_{T-1} + Q_T(h_T, \pi_T(H_T)) | H_{T-1} = h_{T-1}, A_{T-1} = a_{T-1}], \quad (2)$$

and, for $t = T - 2, \dots, 1$,

$$Q_t^{\pi_{(t+1):T}}(h_t, a_t) \equiv E[Y_t + Q_{t+1}^{\pi_{(t+2):T}}(H_{t+1}, \pi_{t+1}(H_{t+1})) | H_t = h_t, A_t = a_t]. \quad (3)$$

We refer to $Q_t^{\pi_{(t+1):T}}(h_t, a_t)$ as the Q-function for $\pi_{(t+1):T}$. The Q-function $Q_t^{\pi_{(t+1):T}}(h_t, a_t)$ represents the average total outcome when the history H_t and treatment A_t correspond to h_t and a_t in stage t , and the future treatment assignments follow $\pi_{(t+1):T}$. When $t = T$, we denote $Q_T^{\pi_{(T+1):T}}(\cdot, \cdot) = Q_T(\cdot, \cdot)$. Note that $E[Q_t^{\pi_{(t+1):T}}(H_t, \pi_t(H_t))] = V_t(\pi_{t:T})$ holds under Assumption 2.1.¹⁰ In what follows, for any function $f(\cdot, \cdot) : \mathcal{H}_t \times \mathcal{A}_t \rightarrow \mathbb{R}$ and policy $\pi_t(\cdot) : \mathcal{H}_t \rightarrow \mathcal{A}_t$, we denote $f(h_t, \pi_t(h_t))$ shortly by $f(h_t, \pi_t)$ (e.g., $Q_t^{\pi_{(t+1):T}}(h_t, \pi_t(h_t))$ is

⁹For any object $v_{s:t}$ and $w_{s:t}$ ($s \leq t$), $v_{t:t}$ and $w_{t:t}$ correspond to v_t and w_t , respectively.

¹⁰See, for example, Tsiatis et al.(2019, Section 6.4) or Lemma B.2 in Appendix B.

denoted by $Q_t^{\pi^{(t+1):T}}(h_t, \pi_t)$.

Given a fixed DTR π , we can use the sequential definitions in equations (1)–(3) to estimate the sequence $\{Q_t^{\pi^{(t+1):T}}(\cdot, \cdot)\}_{t=1, \dots, T}$ of the Q-functions for π . This approach is referred to as the fitted Q-evaluation (Munos and Szepesvári, 2008; Fonteneau et al., 2013; Le et al., 2019) in the reinforcement learning literature and comprises multiple steps as follows:

- Regress Y_T on (H_T, A_T) to obtain $\hat{Q}_T(\cdot, \cdot)$ as the estimated regression function for $Q_T(\cdot, \cdot)$;
- Regress $Y_{T-1} + \hat{Q}_T(H_T, \pi_T)$ on (H_{T-1}, A_{T-1}) to obtain $\hat{Q}_{T-1}^{\pi_T}(\cdot, \cdot)$ as the estimated regression function for $Q_{T-1}^{\pi_T}(\cdot, \cdot)$;
- Recursively, for $t = T - 2, \dots, 1$, regress $Y_t + \hat{Q}_{t+1}^{\pi^{(t+2):T}}(H_{t+1}, \pi_{t+1})$ on (H_t, A_t) to obtain $\hat{Q}_t^{\pi^{(t+1):T}}(\cdot, \cdot)$ as the estimated regression function for $Q_t^{\pi^{(t+1):T}}(\cdot, \cdot)$.

We can apply flexible/nonparametric regression methods, including machine learning methods (e.g., random forests, lasso, neural networks), to the regression in each step.

Given the definitions of the Q-functions, we can optimize the DTR through backward induction. To present the idea, we here assume that the generative distribution function P is known and that the pair (P, Π) satisfies Assumptions 2.1, 2.2, and 2.3. The backward-induction approach in the population problem is a step-wise process that proceeds as follows. In the first step, the approach optimizes the final-stage policy over Π_T by solving

$$\pi_T^{*,B} = \arg \max_{\pi_T \in \Pi_T} E [Q_T(H_T, \pi_T)].$$

Then, recursively, from $t = T - 1$ to 1, the approach optimizes the t -th stage policy over Π_t by solving

$$\pi_t^{*,B} = \arg \max_{\pi_t \in \Pi_t} E \left[Q_t^{\pi^{(t+1):T}}(H_t, \pi_t) \right].$$

Note that the objective function $E \left[Q_t^{\pi^{(t+1):T}}(H_t, \pi_t) \right]$ corresponds to the policy value $V_t(\pi_t, \pi_{(t+1):T}^{*,B})$ under Assumption 2.1. The entire procedure yields the DTR $\pi^{*,B} \equiv (\pi_1^{*,B}, \dots, \pi_T^{*,B})$.

We denote the optimal DTR over Π by $\pi^{*,opt} = \arg \max_{\pi \in \Pi} W(\pi)$. Importantly, the DTR $\pi^{*,B}$ obtained through backward induction does not necessarily correspond to the optimal

DTR $\pi^{*,opt}$ when Π is structurally constrained, as noted by [Li et al. \(2023\)](#) and [Sakaguchi \(2024\)](#). To ensure that backward induction yields the optimal DTR, the policy class Π_t for each stage t (≥ 2) needs to be correctly specified. The following is a sufficient condition for backward optimization to attain optimality.

Assumption 3.1 (Correct Specification). *There exists $\pi_{2:T}^* = (\pi_2^*, \dots, \pi_T^*) \in \Pi_{2:T}$ such that for any $t = 2, \dots, T$,*

$$Q_t^{\pi_{(t+1):T}^*}(H_t, \pi_t^*) \geq \sup_{\pi_t \in \Pi_t} Q_t^{\pi_{(t+1):T}^*}(H_t, \pi_t) \text{ a.s.}$$

Assumption 3.1 requires that, for each stage $t \geq 2$, Π_t includes a policy that can select the optimal arm among all feasible arms in $\{\pi_t(h_t) : \pi_t \in \Pi_t\}$ ($\subseteq \mathcal{A}_t$) for any history h_t . This assumption is satisfied when Π_t (for $t = 2, \dots, T$) is flexible enough or correctly specified in relation to treatment effect heterogeneity.¹¹ We suppose that Π satisfies Assumption 3.1. Note that Assumption 3.1 is a sufficient but not a necessary condition.¹²

A stronger version of Assumption 3.1 is that each Π_t ($t \geq 2$) contains the first-best policy; that is, there exists $\pi_{2:T}^{*,FB} = (\pi_2^{*,FB}, \dots, \pi_T^{*,FB}) \in \Pi_{2:T}$ such that for any $t = 2, \dots, T$,

$$Q_t^{\pi_{(t+1):T}^{*,FB}}(H_t, \pi_t^{*,FB}) \geq \max_{a_t \in \mathcal{A}_t} Q_t^{\pi_{(t+1):T}^{*,FB}}(H_t, a_t) \text{ a.s.}$$

The first-best policy $\pi_t^{*,FB}$ is the policy that selects the best treatment arm for any history h_t . [Li et al. \(2023\)](#) and [Sakaguchi \(2024\)](#) argue that the availability of the first-best policies is a fundamental assumption for the optimality of the backward optimization. However, Assumption 3.1 is practically weaker than this. For example, in the optimal starting problem (Example 2.1), although Π_t may not include the first-best policy $\pi_t^{*,FB}$ due to structural constraints from the optimal starting problem, Assumption 3.1 does not require its feasibility. Similarly, when each policy in Π_t depends only on a sub-history \tilde{h}_t (a subset of h_t), Assumption 3.1 requires that the optimal policy is available for the sub-history, rather than for the entire history.

¹¹[Zhang et al. \(2013\)](#) discuss how to correctly specify Π_t based on the model for treatment effect heterogeneity. [Zhao et al. \(2015\)](#) use a reproducing kernel Hilbert space for each Π_t as a flexible class of policies.

¹²[Sakaguchi \(2024, Remark 3.1\)](#) provides two examples demonstrating cases where backward induction leads to optimality in one instance and sub-optimality in another when Assumption 3.1 does not hold.

The following lemma formalizes the optimality of the backward-induction procedure under Assumption 3.1.

Lemma 3.1. *Under Assumptions 2.1, 2.3, and 3.1, $\pi^{*,B}$ is the optimal DTR over Π ; that is,*

$$W(\pi^{*,B}) \geq W(\pi) \text{ for all } \pi \in \Pi.$$

Proof. See Appendix B.1. □

3.2 Learning of the Optimal DTRs through Backward Induction

This section presents a backward optimization procedure to learn the optimal DTRs using an AIPW estimator of the policy value function. Following the doubly robust policy learning of Athey and Wager (2021) and Zhou et al. (2023b), we employ cross-fitting (Schick, 1986; Chernozhukov et al., 2018) to estimate the policy value function and learn the optimal policy independently. We randomly divide the dataset $\{Z_i : i = 1, \dots, n\}$ into K evenly-sized folds (e.g., $K = 5$). Let I_k be the set of indices of the data in the k -th fold, and I_{-k} denote the set of indices of the data excluded from the k -th fold. In what follows, for any statistics \hat{f} , we denote by \hat{f}^{-k} the corresponding statistics calculated using the data excluded from the k -th fold.

The approach we propose is based on backward induction and thus consists of multiple steps. As a preliminary step, we estimate the propensity scores $\{e_t(\cdot, \cdot)\}_{t=1, \dots, T}$ for all stages and the Q-function $Q_T(\cdot, \cdot)$ for the final stage using the data excluded in each cross-fitting fold. For each fold index k , we denote by $\hat{e}_t^{-k}(\cdot, \cdot)$ and $\hat{Q}_T^{-k}(\cdot, \cdot)$, respectively, the estimators of $e_t(\cdot, \cdot)$ and $Q_T(\cdot, \cdot)$ using data not contained in the k -th fold. Any regression methods, including machine learning methods (e.g., random forests and neural networks), can be used to estimate $e_t(\cdot, \cdot)$ and $Q_T(\cdot, \cdot)$.

Given $\{\hat{e}_t^{-k}(\cdot, \cdot)\}_{t=1, \dots, T}$ and $\hat{Q}_T^{-k}(\cdot, \cdot)$ for each $k = 1, \dots, K$, we estimate the optimal DTR sequentially as follows. In the first step, regarding the final stage T , we construct a score function of the treatment a_T for stage T as

$$\hat{\Gamma}_{i,T}(a_T) \equiv \frac{Y_{i,T} - \hat{Q}_T^{-k(i)}(H_{i,T}, A_{i,T})}{\hat{e}_T^{-k(i)}(H_{i,T}, A_{i,T})} \cdot 1\{A_{i,T} = a_T\} + \hat{Q}_T^{-k(i)}(H_{i,T}, a_T).$$

Given a policy π_T , the sample mean $(1/n) \sum_{i=1}^n \widehat{\Gamma}_{i,T}(\pi_T(H_{i,T}))$ is an AIPW estimator of the policy value $V_t(\pi_T)$ for stage T .

We then find the best candidate policy for stage T by solving

$$\widehat{\pi}_T = \arg \max_{\pi_T \in \Pi_T} \frac{1}{n} \sum_{i=1}^n \widehat{\Gamma}_{i,T}(\pi_T(H_{i,T})).$$

In the next step, we consider stage $T - 1$. Given the estimated policy $\widehat{\pi}_T$ from the previous step, for each k -th cross-fitting fold, we estimate the Q-function $Q_{T-1}^{\widehat{\pi}_T}(\cdot, \cdot)$ for $\widehat{\pi}_T$ by regressing $Y_{i,T-1} + \widehat{Q}_T(H_{i,t}, \widehat{\pi}_T)$ on $(H_{i,T-1}, A_{i,T-1})$ using the observations whose indices are not included in I_k . This corresponds to the second step of the fitted-value Q-evaluation. We denote by $\widehat{Q}_{T-1}^{\widehat{\pi}_T, -k(i)}(\cdot, \cdot)$ the resulting estimator of $Q_{T-1}^{\widehat{\pi}_T}(\cdot, \cdot)$ for each fold k . Any regression method can be applied at this step.

We subsequently construct the score function of a_{T-1} as follows:

$$\begin{aligned} \widehat{\Gamma}_{i,T-1}^{\widehat{\pi}_T}(a_{T-1}) &\equiv \frac{Y_{i,T-1} + \widehat{\Gamma}_{i,T}(\widehat{\pi}_T(H_{i,T})) - \widehat{Q}_{T-1}^{\widehat{\pi}_T, -k(i)}(H_{i,T-1}, A_{i,T-1})}{\widehat{e}_{T-1}^{-k(i)}(H_{i,T-1}, A_{i,T-1})} \cdot \mathbf{1}\{A_{i,T-1} = a_{T-1}\} \\ &\quad + \widehat{Q}_{T-1}^{\widehat{\pi}_T, -k(i)}(H_{i,T-1}, a_{T-1}). \end{aligned}$$

Given a policy π_{T-1} , the sample mean $(1/n) \sum_{i=1}^n \widehat{\Gamma}_{i,T-1}^{\widehat{\pi}_T}(\pi_{T-1}(H_{i,T-1}))$ is an AIPW estimator of the policy value $V_{T-1}(\pi_{T-1}, \widehat{\pi}_T)$. We then find the best candidate policy for stage $T - 1$ by solving

$$\widehat{\pi}_{T-1} = \arg \max_{\pi_{T-1} \in \Pi_{T-1}} \frac{1}{n} \sum_{i=1}^n \widehat{\Gamma}_{i,T-1}^{\widehat{\pi}_T}(\pi_{T-1}(H_{i,T-1})).$$

Recursively, for $t = T - 2, \dots, 1$, we learn the optimal policy as follows. For each cross-fitting index k , we first estimate the Q-function $Q_t^{\widehat{\pi}^{(t+1):T}}$ by regressing $Y_{i,t} + \widehat{Q}_{t+1}^{\widehat{\pi}^{(t+2):T}}(H_{i,t+1}, \widehat{\pi}_{t+1})$ on $(H_{i,t}, A_{i,t})$ using the observations whose indices are not in I_k (the fitted Q-evaluation). Any regression method can be applied at this step.

We next construct the score function of a_t as

$$\begin{aligned} \widehat{\Gamma}_{i,t}^{\widehat{\pi}^{(t+1):T}}(a_t) &\equiv \frac{Y_{i,t} + \widehat{\Gamma}_{i,t+1}^{\widehat{\pi}^{(t+2):T}}(\widehat{\pi}_{t+1}(H_{i,t+1})) - \widehat{Q}_t^{\widehat{\pi}^{(t+1):T}, -k(i)}(H_{i,t}, A_{i,t})}{\widehat{e}_t^{-k(i)}(H_{i,t}, A_{i,t})} \cdot \mathbf{1}\{A_{i,t} = a_t\} \\ &\quad + \widehat{Q}_t^{\widehat{\pi}^{(t+1):T}, -k(i)}(H_{i,t}, a_t). \end{aligned}$$

We then find the best candidate policy for stage t by solving

$$\hat{\pi}_t = \arg \max_{\pi_t \in \Pi_t} \frac{1}{n} \sum_{i=1}^n \hat{\Gamma}_{i,t}^{\hat{\pi}_{(t+1):T}}(\pi_t(H_{i,t})),$$

where the objective function $(1/n) \sum_{i=1}^n \hat{\Gamma}_{i,t}^{\hat{\pi}_{(t+1):T}}(\pi_t(H_{i,t}))$ is an AIPW estimator of the policy value $V_t(\pi_t, \hat{\pi}_{(t+1):T})$.

Throughout this procedure, we obtain the sequence $\hat{\pi} \equiv (\hat{\pi}_1, \dots, \hat{\pi}_T)$, which serves as the estimator for the optimal DTR. In the following section, we will show its statistical properties with respect to its welfare regret.

4 Statistical Properties

Given a DTR $\pi \in \Pi$, we define the regret of π by $R(\pi) \equiv \sup_{\tilde{\pi} \in \Pi} W(\tilde{\pi}) - W(\pi)$, the loss of the welfare of π relative to the maximum welfare achievable in Π . We study statistical properties of $\hat{\pi}$ with respect to its regret $R(\hat{\pi})$. This section shows the rate of convergence of $R(\hat{\pi})$ depending on the convergence rates of the estimators of the nuisance components, $\{\hat{e}_t^{-k}(\cdot, \cdot)\}_{t=1, \dots, T}$ and $\{\hat{Q}_t^{\pi_{(t+1):T}, -k}(\cdot, \cdot)\}_{t=1, \dots, T}$, and the complexity of Π .

Let $\hat{Q}_t^{\pi_{(t+1):T}, (n)}(\cdot, \cdot)$ and $\hat{e}_t^{(n)}(\cdot, \cdot)$ denote the estimators of the Q-function $Q_t^{\pi_{(t+1):T}}(\cdot, \cdot)$ for $\pi_{(t+1):T}$ and the propensity score $e_t(\cdot, \cdot)$, respectively, using a sample of size n randomly drawn from the distribution P . For $t = T$, we denote $\hat{Q}_T^{\pi_{(T+1):T}, (n)}(\cdot, \cdot) = \hat{Q}_T^{(n)}(\cdot, \cdot)$. We suppose that $\{\hat{Q}_t^{\pi_{(t+1):T}, (n)}(\cdot, \cdot)\}_{t=1, \dots, T}$ and $\{\hat{e}_t^{(n)}(\cdot, \cdot)\}_{t=1, \dots, T}$ satisfy the following assumption.

Assumption 4.1. (i) *There exists $\tau > 0$ such that the following holds: For all $t = 1, \dots, T$, $s = 1, \dots, t$, and $m \in \{0, 1\}$,*

$$\begin{aligned} & \sup_{\underline{a}_{s:t} \in \underline{\mathcal{A}}_{s:t}} E \left[\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \left(\hat{Q}_t^{\pi_{(t+1):T}, (n)}(H_t, a_t) - Q_t^{\pi_{(t+1):T}}(H_t, a_t) \right)^2 \right] \\ & \times E \left[\left(\frac{1}{\prod_{\ell=s}^{t-m} \hat{e}_\ell^{(n)}(H_\ell, a_\ell)} - \frac{1}{\prod_{\ell=s}^{t-m} e_\ell(H_\ell, a_\ell)} \right)^2 \right] = \frac{o(1)}{n^\tau}. \end{aligned}$$

(ii) *There exists $n_0 \in \mathbb{N}$ such that for any $n \geq n_0$ and $t = 1, \dots, T$,*

$$\sup_{a_t \in \mathcal{A}_t, \pi_{(t+1):T} \in \Pi_{(t+1):T}} \hat{Q}_t^{\pi_{(t+1):T}, (n)}(H_t, a_t) < \infty \quad \text{and} \quad \min_{a_t \in \mathcal{A}_t} \hat{e}_t^{(n)}(H_t, a_t) > 0$$

hold *a.s.*

As we will see later, the \sqrt{n} -consistency of the regret $R(\hat{\pi})$ to zero can be achieved when Assumption 4.1 (i) holds with $\tau = 1$, which is not very strong or restrictive. For example, Assumption 4.1 (i) is satisfied with $\tau = 1$ when

$$\begin{aligned} & \sup_{a_t \in \mathcal{A}_t} E \left[\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \left(\widehat{Q}_t^{\pi_{(t+1):T}, (n)}(H_t, a_t) - Q_t^{\pi_{(t+1):T}}(H_t, a_t) \right)^2 \right] = \frac{o(1)}{\sqrt{n}} \text{ and} \\ & \sup_{a_{s:t} \in \mathcal{A}_{s:t}} E \left[\left(\frac{1}{\prod_{\ell=s}^t \widehat{e}_\ell^{(n)}(H_\ell, a_\ell)} - \frac{1}{\prod_{\ell=s}^t e_\ell(H_\ell, a_\ell)} \right)^2 \right] = \frac{o(1)}{\sqrt{n}} \end{aligned}$$

hold for all $t = 1, \dots, T$ and $s = 1, \dots, t$. The uniform MSE convergence rate of $\widehat{Q}_t^{\pi_{(t+1):T}, (n)}(H_t, a_t)$ for the fitted Q-evaluation is not a standard result, but some existing results could be applicable with some modifications (e.g., Zhang et al., 2018).¹³ Note also that Assumption 4.1 (i) encompasses the property of doubly robustness; that is, Assumption 4.1 (i) holds if either $\widehat{Q}_t^{\pi_{(t+1):T}, (n)}(\cdot, \cdot)$ is uniformly consistent or $\widehat{e}_t^{(n)}(\cdot, \cdot)$ is consistent.

We next consider the complexity of the class Π of DTRs and the class Π_t of stage-specific policies. Following Zhou et al. (2023b), we use the ϵ -Hamming covering number to measure the complexity of the class of policy sequences $\pi_{s:t} \in \Pi_s \times \dots \times \Pi_t$ for each s and t such that $s \leq t$.

Definition 4.1. (i) For any stages s and t such that $s \leq t$, given a set of history points $\{h_t^{(1)}, \dots, h_t^{(n)}\} \subseteq \mathcal{H}_t$, we define the Hamming distance between two sequences of policies $\pi_{s:t}, \pi'_{s:t} \in \Pi_{s:t}$ as $d_h(\pi_{s:t}, \pi'_{s:t}) \equiv n^{-1} \sum_{i=1}^n 1\{\pi_s(h_s^{(i)}) \neq \pi'_s(h_s^{(i)}) \vee \dots \vee \pi_t(h_t^{(i)}) \neq \pi'_t(h_t^{(i)})\}$, where $h_s^{(i)} \subseteq h_{s+1}^{(i)} \subseteq \dots \subseteq h_t^{(i)} \in \mathcal{H}_t$ from the definition of the history. Let $N_{d_h}(\epsilon, \Pi_{s:t}, \{h_t^{(1)}, \dots, h_t^{(n)}\})$ be the smallest number of sequences of policies $\pi_{s:t}^{(1)}, \pi_{s:t}^{(2)}, \dots$ in $\Pi_{s:t}$ such that for any $\pi_{s:t} \in \Pi_{s:t}$, there exists $\pi_{s:t}^{(i)}$ satisfying $d_h(\pi_{s:t}, \pi_{s:t}^{(i)}) \leq \epsilon$. We define the ϵ -Hamming covering number of $\Pi_{s:t}$ as

$$N_{d_h}(\epsilon, \Pi_{s:t}) \equiv \sup \left\{ N_{d_h} \left(\epsilon, \Pi_{s:t}, \{h_t^{(1)}, \dots, h_t^{(n)}\} \right) \mid n \geq 1, h_t^{(1)}, \dots, h_t^{(n)} \in \mathcal{H}_t \right\}.$$

(ii) We define the entropy integral of $\Pi_{s:t}$ as $\kappa(\Pi_{s:t}) = \int_0^1 \sqrt{\log N_{d_h}(\epsilon^2, \Pi_{s:t})} d\epsilon$.

¹³Zhang et al. (2018) derive the rate of convergence for offline Q-learning using Support Vector Machine regression with policy search over the class of list-form policies. While their study focuses on Q-learning, the results can likely be extended to fitted Q-evaluation.

Note that when $s = t$, $N_{d_h}(\epsilon, \Pi_{s:t}) = N_{d_h}(\epsilon, \Pi_t)$ and $\kappa(\Pi_{s:t}) = \kappa(\Pi_t)$. We assume that the policy class Π_t for each t is not excessively complex in terms of the covering number.

Assumption 4.2. *For all $t = 1, \dots, T$, $N_{d_h}(\epsilon, \Pi_t) \leq C \exp(D(1/\epsilon)^\omega)$ holds for any $\epsilon > 0$ and some constants $C, D > 0$ and $0 < \omega < 0.5$.*

This assumption implies that the covering number of Π_t does not grow too quickly, but allows $\log N_{d_h}(\epsilon, \Pi_t)$ to grow at a rate of $1/\epsilon$. This assumption is satisfied, for example, by a class of finite-depth trees (see Zhou et al. (2023b, Lemma 4)). In the case of a binary action set (i.e., $|\mathcal{A}_t| = 2$), a VC-class of \mathcal{A}_t also satisfies Assumption 4.2. Zhou et al. (2023b, Remark 4) shows that the entropy integral $\kappa(\Pi_t)$ is finite under Assumption 4.2.

Regarding the class Π of entire DTRs, its entropy integral $\kappa(\Pi)$ is finite as well under Assumption 4.2.

Lemma 4.2. *Under Assumption 4.2, $\kappa(\Pi) < \infty$.*

Proof. See Appendix B.1. □

The following theorem presents the main result of this paper, showing the rate of convergence for the regret of the DTR $\hat{\pi}$ obtained using the proposed approach.

Theorem 4.3. *Under Assumptions 2.1–2.3, 3.1, 4.1, and 4.2,*

$$R(\hat{\pi}) = O_p(\kappa(\Pi) \cdot n^{-1/2}) + O_p(n^{-\min\{1/2, \tau/2\}}).$$

Proof. See Appendix A. □

This theorem shows the convergence rate of the regret $R(\hat{\pi})$ for the proposed approach. When Assumption 4.1 (i) holds with $\tau = 1$, the approach achieves the minimax optimal rate $\kappa(\Pi) \cdot n^{-1/2}$ for the regret convergence.¹⁴ This result is comparable to those of Athey and Wager (2021) and Zhou et al. (2023b), who study doubly robust policy learning in single-stage settings.¹⁵

¹⁴In the case of binary treatment assignment at each stage, Sakaguchi (2024) shows that the minimax optimal rate of convergence for the regret is $V_{1:T} \cdot n^{-1/2}$, where $V_{1:T}$ is the VC-dimension of the class of DTRs.

¹⁵In the single-stage binary treatment setting, Athey and Wager (2021) use the specific growth rate of the entropy in a VC class and obtain a slightly stronger result compared to using the fixed entropy class.

In the proof of Theorem 4.3, we consider the derivation of the asymptotic upper bound on $R(\hat{\pi})$. This is, however, a nontrivial task because the stage-specific policies in $\hat{\pi} = (\hat{\pi}_1, \dots, \hat{\pi}_T)$ are estimated sequentially, rather than simultaneously. If the DTR is estimated simultaneously across all stages, we can apply the theoretical analysis of Zhou et al. (2023b) with some modifications. However, since $\hat{\pi}$ is estimated sequentially, we cannot directly apply their approach. While the sequential estimation complicates the analysis of $R(\hat{\pi})$, Appendix A introduces new analytical strategies for evaluating the regret of the sequentially estimated DTR.

5 Existing Approach

An alternative doubly robust approach for estimating the optimal DTR is maximizing an AIPW estimator of the welfare function over the entire class of DTRs. Specifically, we have

$$\begin{aligned} \hat{\pi}^{AIPW} &= \arg \max_{\pi \in \Pi} \widehat{W}^{AIPW}(\pi) \quad \text{with} \\ \widehat{W}^{AIPW}(\pi) &= \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \left(\hat{\psi}_{it}^{-k(i)}(\underline{\pi}_t) \cdot Y_{it} \right. \\ &\quad \left. - \left(\hat{\psi}_{it}^{-k(i)}(\underline{\pi}_t) - \hat{\psi}_{i,t-1}^{-k(i)}(\underline{\pi}_{t-1}) \right) \cdot \widehat{Q}_t^{\pi(t+1):T, -k(i)}(H_{it}, \pi_t(H_{it})) \right) \end{aligned} \quad (4)$$

and $\hat{\psi}_{it}^{-k(i)}(\underline{\pi}_t) \equiv \left(\prod_{s=1}^t 1 \{A_{is} = \pi_s(H_{is})\} \right) / \left(\prod_{s=1}^t \hat{e}_t^{-k(i)}(H_{is}, A_{is}) \right)$, where $\widehat{W}^{AIPW}(\pi)$ is an AIPW estimator of $W(\pi)$. This approach is first proposed by Zhang et al. (2013) without using cross-fitting, but they did not show its statistical properties.¹⁶ One advantage of this method is that it does not require the correct specification of Π (Assumption 3.1) for consistent estimation of the optimal DTR.

However, this approach faces two computational challenges. First, since the nuisance components $\{Q_t^{\pi(t+1):T}\}_{t=1, \dots, T}$ depend on each specific DTR π , we basically need to estimate $\{Q_t^{\pi(t+1):T}\}_{t=1, \dots, T}$ for every DTR π to implement this method. Nie et al. (2021) highlight this computational issue. Second, maximizing $\widehat{W}^{AIPW}(\pi)$ simultaneously across all stages is computationally difficult, particularly when T is not very small. These computational issues render the approach intractable unless the class of DTRs is small (e.g., Π consisting of a small number of DTRs). By contrast, our proposed approach has two key

¹⁶Jiang and Li (2016), Thomas and Brunskill (2016), and Kallus and Uehara (2020) propose AIPW estimators of welfare functions for evaluating fixed DTRs, but their focus is not on optimizing DTRs.

advantages: (i) the nuisance components $\{Q_t^{\hat{\pi}^{(t+1):T}}\}_{t=1,\dots,T}$ depend only on the estimated policies $\hat{\pi}_{(t+1)}$ from the earlier steps of the backward optimization, and (ii) the backward optimization is computationally much more efficient than the simultaneous optimization across all stages.

Though the simultaneous maximization approach (4) is not our primary proposal, we show its statistical properties as follows.

Theorem 5.1. *Under Assumptions 2.1–2.3, 4.1, and 4.2,*

$$R(\hat{\pi}^{AIPW}) = O_p(\kappa(\Pi) \cdot n^{-1/2}) + O_p(n^{-\min\{1/2, \tau/2\}}). \quad (5)$$

Proof. See Appendix C. □

This theorem shows that $\hat{\pi}^{AIPW}$ attains the same convergence rate for regret as our proposed sequentially estimated DTR, $\hat{\pi}$, under the same conditions regarding the MSE convergence rate of the nuisance component estimators and the complexity of the DTR class. Since the simultaneous optimization approach does not require the correct specification of the DTR class, Theorem 5.1 does not rely on Assumption 3.1.

6 Simulation Study

We conduct a simulation study to examine the finite sample performance of the approach presented in Section 3. We consider two data generating processes (DGPs), labeled DGP1 and DGP2, each of which consists of two stages of binary treatment assignment $(A_1, A_2) \in \{0, 1\}^2$, associated potential outcomes $\{Y_2(a_1, a_2)\}_{(a_1, a_2) \in \{0, 1\}^2}$ for the second stage, 20 state variables $(S_1^{(1)}, \dots, S_1^{(20)})$ observed at the first stage, and one state variable S_2 observed at the second stage. Each DGP is structured as follows:

$$(S_1^{(1)}, \dots, S_1^{(20)})' \sim N(\mathbf{0}, I_{20});$$

$$S_2(a_1) = \text{sign}\left(S_1^{(1)}\right) \cdot a_1 + S_1^{(2)} + \left(S_1^{(3)}\right)^2 + S_1^{(4)} + \varepsilon_1 \text{ with } \varepsilon_1 \sim N(0, 1);$$

$$Y_2(a_1, a_2) = \phi(a_1, S_2(a_1)) \cdot a_2 + 0.5 \cdot S_2(a_1) + S_1^{(4)} - \left(S_1^{(5)}\right)^2 + S_1^{(6)} + \varepsilon_2 \text{ with } \varepsilon_2 \sim N(0, 1);$$

$$A_1 \sim \text{Ber}\left(1/(1 + e^{0.5S_1^{(2)} - 0.5S_1^{(3)} - S_1^{(5)}})\right), \quad A_2 \sim \text{Ber}\left(1/(1 + e^{0.5S_1^{(5)} + 0.5S_2 - 0.2A_1})\right).$$

In DGPs 1 and 2, we specify $\phi(a_1, S_2(a_1))$ as

$$\begin{aligned}\phi(a_1, S_2(a_1)) &= \text{sign}(S_2(a_1) \cdot (a_1 - 1/2)) \quad \text{and} \\ \phi(a_1, S_2(a_1)) &= S_2(a_1) + (a_1 - 1/2),\end{aligned}$$

respectively. In each DGP, the first-stage treatment a_1 influences the outcome Y_2 through both direct and indirect channels: (i) a direct effect on Y_2 via treatment effect heterogeneity $Y_2(a_1, 1) - Y_2(a_1, 0)$ and (ii) an indirect effect on Y_2 via the second-stage state variable $S_2(a_1)$.

We compare the performance of the approach proposed in Section 3 (labeled “DR”) with those of three existing methods: Q-learning without and with policy search (labeled “Q-learn” and “Q-search,” respectively) and the IPW classification-based approach with backward optimization (labeled “IPW”).¹⁷ For each method, we use generalized random forests (Athey et al., 2019) to estimate nuisance components. We set $K = 5$ for cross-fitting in the proposed approach.

For DR, Q-search, and IPW, we use a class of DTRs $\Pi = \Pi_1 \times \Pi_2$ with Π_1 being the class of depth-1 decision trees of H_1 , and Π_2 being the class of depth-2 decision trees of H_2 . In DGP1, Π_2 is correctly specified in the sense of Assumption 3.1, whereas in DGP2, Π_2 is misspecified, leading to a loss of optimality in backward optimization. Note that Q-learn consistently estimates the optimal DTR in both DGPs. We solve the optimization problems involving decision trees using the exact learning algorithm for decision trees proposed by Zhou et al. (2023b).

[Table 1 about here]

Table 1 presents the results of 500 simulations with sample sizes of $n = 250, 500, 1000, 2000,$ and 4000 . In each simulation, welfare is calculated using a test sample of 50,000 observations randomly drawn from the same DGP. The results show that DR consistently outperforms the other methods for both DGPs and all sample sizes in terms of mean

¹⁷Following Murphy (2005), we separately consider Q-learning with and without policy search. In Q-learning with policy search, the optimal policy for each stage t is chosen from a pre-specified policy class Π_t ; specifically, the optimal policy is estimated as $\hat{\pi}_t = \arg \max_{\pi_t \in \Pi_t} \sum_{i=1}^n \hat{Q}_t^{\hat{\pi}_t(t+1):T}(H_{it}, \pi_t)$. This approach, for example, is used by Zhang et al. (2018) with a class of list-form policies for Π_t . In contrast, Q-learning without policy search optimizes the policy for each stage t over all measurable policies, such as $\hat{\pi}_t(h_t) = \arg \max_{a_t \in \mathcal{A}_t} \sum_{i=1}^n \hat{Q}_t^{\hat{\pi}_t(t+1):T}(h_t, a_t)$ for any h_t . This approach consistently estimates the first-best DTR unless the Q-functions are not misspecified.

welfare. For example, in DGP1, with a small sample size ($n = 500$), DR achieves over 40% higher welfare than all other methods. Notably, DR surpasses Q-learn even in DGP2, where DR does not consistently estimate the optimal DTR, but Q-learn does.

7 Empirical Application

We apply the proposed approach to data from Project STAR (e.g., [Krueger, 1999](#); [Ding and Lehrer, 2010](#)), where we study the optimal allocation of students to regular-size classes (22 to 25 students per teacher) with a full-time teacher aide and small-size classes (13 to 17 students per teacher) without one in their early education (kindergarten and grade 1).¹⁸ We use a dataset of 1,877 students who were not assigned to regular-size classes without a full-time teacher aide in kindergarten.¹⁹ Among these students, 702 were randomly assigned to regular-size classes with a teacher aide, whereas the remaining students were assigned to small-size classes without a teacher aide in kindergarten (labeled “grade K”). Upon their progression to grade 1, the students were intended to be randomly reassigned between the two class types. However, as some students selected the class types themselves, the allocation process was not entirely random (see, e.g., [Ding and Lehrer \(2010\)](#) for a detailed discussion). Therefore, we consider this empirical task in the observational data setting.

We investigate the optimal allocation of students to the two class types in grades K and 1, based on their socioeconomic backgrounds, educational environment, and intermediate academic achievement. Each student’s academic achievement is measured by their percentile rank on combined reading and mathematics test scores taken at the end of grade 1. The welfare objective for the policymaker is assumed to be maximizing the population average of this academic achievement measure.

We define the first and second stages ($t = 1$ and $t = 2$) as grades K and 1, respectively. We define the action set $\mathcal{A}_t = \{\text{aide, small}\}$, where the treatment variable A_t is labeled “aide” if a student is in a regular-size class with a teacher aide at stage t , and “small” if the student is in a small-size class. The outcome variable Y_2 denotes the percentile rank of the combined reading and mathematics scores at the end of grade 1. We do not

¹⁸[Krueger \(1999\)](#) reports that the presence of a teacher aide did not have a significant impact on student test scores. However, whether teaching aides have effects on academic attainment has not been examined by accounting for multiple stages of treatment and treatment effect heterogeneity.

¹⁹We exclude regular-size classes without a teacher aide, as they are unlikely to be preferable to either regular-size classes with an aide or small-size classes without one for any student.

incorporate the first-stage outcome into the objective function $W(\pi)$.

We include seven variables in $H_1(= S_1)$: student gender, student ethnicity (White/Asian or other), eligibility for free or reduced-price school lunch, school location type (rural or non-rural), teacher’s degree (bachelor’s or higher), years of teaching experience, and teacher ethnicity (White or other). For the second-stage state variables S_2 , we include three variables: reading, math, and total test scores at the end of kindergarten. Recall that $H_2 = (A_1, S_1, S_2)$.

The class of DTRs $\Pi = \Pi_1 \times \Pi_2$ is defined as follows. For the policy class Π_1 associated with kindergarten class allocation, we employ a class of depth-1 trees that may take splitting variables from teacher degree, teacher experience, and school location type. For the policy class Π_2 associated with class allocation in grade 1, we use a class of depth-2 trees that may take splitting variables from reading, math, and total test scores at the end of kindergarten, as well as the kindergarten class type. Note that we exclude student gender, student ethnicity, and teacher ethnicity as splitting variables, as using such variables for treatment choice would be discriminatory.

[Figure 1 about here]

Figure 1 displays the DTR estimated using the proposed approach. The policy for kindergarten class allocation uses teacher experience to determine class type, indicating that teachers with 19 years of experience or less should be assigned to small-size classes. The policy for grade 1 class allocation uses total and reading test scores at the end of kindergarten to determine each student’s class type. For example, following the estimated policy, students with a total test score below 914 are assigned to the small-size class in grade 1.

[Table 2 about here]

Table 2 presents an estimate of the welfare contrast of the optimal DTR $\pi^{*,opt}$ relative to the uniform assignment of all students to small classes in both grades (i.e., $W(\pi^{*,opt}) - E[Y_2(\text{small}, \text{small})]$), along with the share of students in each of the four allocation arms. These estimates are obtained using 5-fold cross-validation. The results show that class-type allocation with the optimal DTR improves academic achievement by an average of 1.27% compared with the uniform allocation of small-size classes. Additionally, under the DTR allocation, around 23% of students are placed in regular-size classes with a teacher

aide in either grade K or 1. Given the higher cost of small-size classes relative to regular-size classes with a teacher aide on a per-student basis (Word et al., 1990)²⁰, this finding suggests that allocation based on the DTR can reduce the costs associated with class-size reduction while simultaneously improving students' academic achievement.

8 Conclusion

We studied the statistical learning of the optimal DTR using observational data and developed a doubly robust approach to learning it under the assumption of sequential ignorability. Based on backward induction, the approach learns the optimal DTR sequentially, ensuring computational tractability. Our main result shows that the DTR obtained through this approach achieves a convergence rate of $n^{-1/2}$ for welfare regret under mild conditions on the MSE convergence rate for estimators of the propensity scores and Q-functions. The simulation study confirms that the proposed approach outperforms others in finite sample settings. Applying the proposed approach to Project STAR data, we estimate the optimal DTR for the sequential allocation of students to regular-size classes with a teacher aide and small-size classes without one in early education.

²⁰According to Word et al. (1990), adding a full-time aide in Grades K-3 across Tennessee cost approximately 75 million dollars annually, while reducing class sizes by one-third cost around 196 to 205 million dollars per year.

Tables

Table 1: Monte Carlo simulation results

Method	DGP1 & Size(n)					DGP2 & Size(n)				
	250	500	1000	2000	4000	250	500	1000	2000	4000
Q-learn	0.20 (0.11)	0.31 (0.11)	0.44 (0.10)	0.58 (0.07)	0.69 (0.03)	1.29 (0.36)	1.69 (0.13)	1.77 (0.07)	1.82 (0.05)	1.87 (0.03)
Q-search	0.17 (0.19)	0.29 (0.22)	0.41 (0.21)	0.54 (0.16)	0.65 (0.10)	1.27 (0.45)	1.60 (0.21)	1.53 (0.10)	1.51 (0.06)	1.53 (0.09)
IPW	0.23 (0.13)	0.30 (0.15)	0.37 (0.16)	0.46 (0.15)	0.55 (0.12)	1.32 (0.2)	1.41 (0.18)	1.52 (0.19)	1.66 (0.21)	1.80 (0.19)
DR	0.28 (0.18)	0.43 (0.17)	0.62 (0.11)	0.70 (0.04)	0.73 (0.02)	1.50 (0.26)	1.76 (0.24)	1.93 (0.14)	1.99 (0.07)	2.02 (0.03)

Notes: Each cell shows the mean welfare, with the standard deviation in parentheses, for each method, DGP, and sample size. These values are calculated based on 500 simulations using a test sample of 50,000 observations randomly drawn from the same DGP. Columns 2-6 and 7-11 show the results for DGP1 and DGP2, respectively, across sample sizes of 250, 500, 1000, 2000, and 4000.

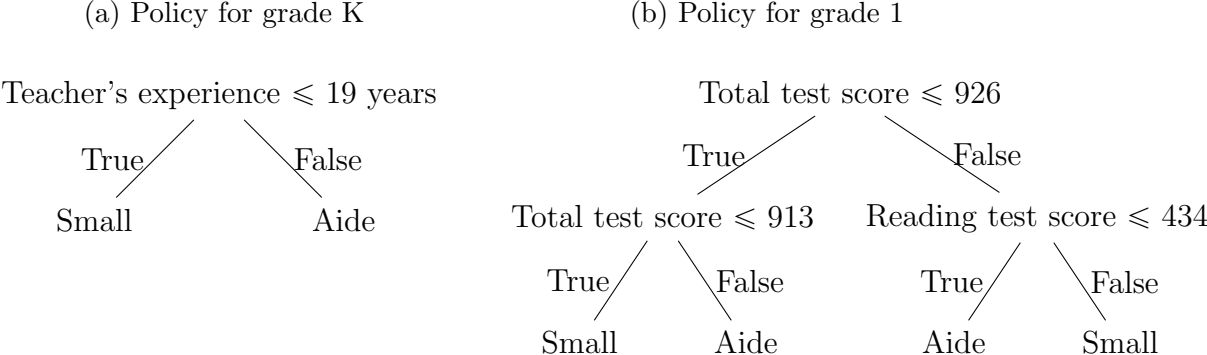
Table 2: Empirical results for optimal class-type allocation

Welfare contrast	Share of students in each allocation arm (a_1, a_2)			
	(aide, aide)	(small, aide)	(aide, small)	(small, small)
1.27%	1.0%	17.2%	5.1%	76.7%

Notes: The first column shows the estimate of the welfare contrast, defined as $W(\pi^{*,opt}) - E[Y_2(\text{small, small})]$. Each of the second to fifth columns shows the estimated share of students assigned to each of the four allocation arms, $(a_1, a_2) \in \{\text{aide, small}\}^2$, by the optimal DTR.

Figure

Figure 1: Estimated DTR for class assignment in grades K and 1



Notes: This figure illustrates the DTR estimated in Section 7. Panels (a) and (b) show the estimated policy trees for grades K and 1, respectively.

Appendix

A Proof of Theorem 4.3

This appendix presents the proof of Theorem 4.3, along with some auxiliary lemmas. We aim to derive an asymptotic upper bound on $R(\hat{\pi})$. This is however a non-trivial task because the components of the DTR, $\hat{\pi} = (\hat{\pi}_1, \dots, \hat{\pi}_T)$, are estimated sequentially rather than simultaneously. We hence cannot directly apply the theoretical analysis of [Athey and Wager \(2021\)](#) and [Zhou et al. \(2023b\)](#), which study doubly robust policy learning in single-stage settings. In the following, we present an original analysis to derive an asymptotic upper bound on $R(\hat{\pi})$.

We begin by noting that Lemma 3.1 allows us to analyze the welfare regret, $R(\hat{\pi}) = W(\pi^{*,opt}) - W(\hat{\pi})$, by evaluating the welfare difference $W(\pi^{*,B}) - W(\hat{\pi})$ between the DTR $\pi^{*,B}$ derived from the backward optimization in the population problem and the estimated DTR $\hat{\pi}$. This approach simplifies the analysis of $R(\hat{\pi})$.

Given the estimated DTR $\hat{\pi}$, we define $R_t^{\hat{\pi}_{t:T}}(\pi_t) \equiv V_t(\pi_t, \hat{\pi}_{(t+1):T}) - V_t(\hat{\pi}_{t:T})$ for any $\pi_t \in \Pi_t$ and $t = 1, \dots, T$. $R_t^{\hat{\pi}_{t:T}}(\pi_t)$ measures the deviation of the policy π_t from the sequence of estimated policies $\hat{\pi}_{t:T}$ at stage t with respect to the value function. We denote $R_T^{\hat{\pi}_{T:T}}(\pi_T) = V_T(\pi_T) - V_T(\hat{\pi}_T)$ for $t = T$.

The following lemma provides a useful result for analyzing the regret $R(\hat{\pi})$, relating the regret of the entire DTR to the stage-specific regrets.

Lemma A.1. *Under Assumptions 2.1, 2.3, and 3.1, the regret of $\hat{\pi}$ is bounded from above as*

$$R(\hat{\pi}) \leq R_1^{\hat{\pi}_{1:T}}(\pi_1^{*,B}) + \sum_{t=2}^T \frac{2^{t-2}}{\eta^{t-1}} R_t^{\hat{\pi}_{t:T}}(\pi_t^{*,B}). \quad (\text{A.1})$$

Proof. See Appendix B.2. □

The result (A.1) enables us to evaluate $R(\hat{\pi})$ through evaluating stage-specific regrets $R_t^{\hat{\pi}_{t:T}}(\pi_t^{*,B})$ ($t = 1, \dots, T$), which is simpler to analyze as we will see.

Given a fixed DTR $\pi = (\pi_1, \dots, \pi_T)$, we define

$$\begin{aligned}\tilde{V}_{i,T}(\pi_T) &\equiv \frac{Y_{i,T} - Q_T(H_{i,T}, A_{i,T})}{e_T(H_{i,T}, A_{i,T})} \cdot \mathbb{1}\{A_{i,T} = \pi_T(H_{i,T})\} + Q_T(H_{i,T}, \pi_T(H_{i,T})), \\ \hat{V}_{i,T}(\pi_T) &\equiv \frac{Y_{i,T} - \hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,T}, A_{i,T})}{\hat{e}_T^{-k(i)}(H_{i,T}, A_{i,T})} \cdot \mathbb{1}\{A_{i,T} = \pi_T(H_{i,T})\} \\ &\quad + \hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,T}, \pi_T(H_{i,T})),\end{aligned}$$

and recursively for $t = T - 1, \dots, 1$:

$$\begin{aligned}\tilde{V}_{i,t}(\pi_{t:T}) &\equiv \frac{Y_{i,t} + \tilde{V}_{i,t+1}(\pi_{(t+1):T}) - Q_t^{\pi_{(t+1):T}}(H_{i,t}, A_{i,t})}{e_t(H_{i,t}, A_{i,t})} \cdot \mathbb{1}\{A_{i,t} = \pi_t(H_{i,t})\} \\ &\quad + Q_t^{\pi_{(t+1):T}}(H_{i,t}, \pi_t(H_{i,t})), \\ \hat{V}_{i,t}(\pi_{t:T}) &\equiv \frac{Y_{i,t} + \hat{V}_{i,t+1}(\pi_{(t+1):T}) - \hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,t}, A_{i,t})}{\hat{e}_t^{-k(i)}(H_{i,t}, A_{i,t})} \cdot \mathbb{1}\{A_{i,t} = \pi_t(H_{i,t})\} \\ &\quad + \hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,t}, \pi_t(H_{i,t})).\end{aligned}$$

Note that the sample mean $(1/n) \sum_{i=1}^n \tilde{V}_{i,t}(\pi_{t:T})$ is an oracle estimate of the policy value function $V_t(\pi_{t:T})$ with oracle access to $\{Q_s^{\pi_{s:T}}(\cdot, \cdot)\}_{s=t+1, \dots, T}$ and $\{e_s(\cdot, \cdot)\}_{s=t, \dots, T}$. Lemma B.2 in Appendix B.1 shows that $(1/n) \sum_{i=1}^n \tilde{V}_{i,t}(\pi_{t:T})$ is an unbiased estimator of the policy value $V_t(\pi_{t:T})$ under the assumption of sequential ignorability (Assumption 2.1).

Following the analysis of Zhou et al. (2023b), for each stage $t = 1, \dots, T$, we define the policy value difference function $\Delta_t(\cdot; \cdot) : \Pi_{t:T} \times \Pi_{t:T} \rightarrow \mathbb{R}$, the oracle influence difference function $\tilde{\Delta}_t(\cdot; \cdot) : \Pi_{t:T} \times \Pi_{t:T} \rightarrow \mathbb{R}$, and the estimated policy value difference function $\hat{\Delta}_t(\cdot; \cdot) : \Pi_{t:T} \times \Pi_{t:T} \rightarrow \mathbb{R}$, as follows: For $\pi_{t:T}^a = (\pi_t^a, \dots, \pi_T^a) \in \Pi_{t:T}$ and $\pi_{t:T}^b = (\pi_t^b, \dots, \pi_T^b) \in \Pi_{t:T}$,

$$\Delta_t(\pi_{t:T}^a; \pi_{t:T}^b) \equiv V_t(\pi_{t:T}^a) - V_t(\pi_{t:T}^b), \quad (\text{A.2})$$

$$\tilde{\Delta}_t(\pi_{t:T}^a; \pi_{t:T}^b) \equiv \frac{1}{n} \sum_{i=1}^n \tilde{V}_{i,t}(\pi_{t:T}^a) - \frac{1}{n} \sum_{i=1}^n \tilde{V}_{i,t}(\pi_{t:T}^b), \quad (\text{A.3})$$

$$\hat{\Delta}_t(\pi_{t:T}^a; \pi_{t:T}^b) \equiv \frac{1}{n} \sum_{i=1}^n \hat{V}_{i,t}(\pi_{t:T}^a) - \frac{1}{n} \sum_{i=1}^n \hat{V}_{i,t}(\pi_{t:T}^b).$$

From the definitions, the stage-specific regret $R_t^{\hat{\pi}_{t:T}}(\pi_t^{*,B})$ is expressed as

$$R_t^{\hat{\pi}_{t:T}}(\pi_t^{*,B}) = \Delta_t\left(\pi_t^{*,B}, \hat{\pi}_{(t+1):T}; \hat{\pi}_{t:T}\right).$$

In what follows, we evaluate $R_t^{\hat{\pi}_{t:T}}(\pi_t^{*,B})$ for each t . A standard argument of the statistical learning theory (e.g., [Lugosi, 2002](#)) gives

$$\begin{aligned}
R_t^{\hat{\pi}_{t:T}}(\pi_t^{*,B}) &= \Delta_t \left(\pi_t^{*,B}, \hat{\pi}_{(t+1):T}; \hat{\pi}_{t:T} \right) \\
&\leq \Delta_t \left(\pi_t^{*,B}, \hat{\pi}_{(t+1):T}; \hat{\pi}_{t:T} \right) - \hat{\Delta}_t \left(\pi_t^{*,B}, \hat{\pi}_{(t+1):T}; \hat{\pi}_{t:T} \right) \\
&\leq \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} \left| \Delta_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \hat{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) \right| \\
&\leq \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} \left| \Delta_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) \right| \\
&+ \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} \left| \hat{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) \right|,
\end{aligned} \tag{A.4}$$

where the first inequality follows because $\hat{\pi}_t$ maximizes $(1/n) \sum_{i=1}^n \hat{\Gamma}_{i,t}^{\hat{\pi}_{(t+1):T}}(\pi_t(H_{i,t}))$ over Π_t ; hence, $\hat{\Delta}_t \left(\pi_t^{*,B}, \hat{\pi}_{(t+1):T}; \hat{\pi}_{t:T} \right) \leq 0$.

We can now evaluate $R_t^{\hat{\pi}_{t:T}}(\pi_t^{*,B})$ through evaluating

$$\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} \left| \Delta_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) \right| \tag{A.5}$$

and

$$\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} \left| \hat{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) \right|. \tag{A.6}$$

As for the former, we apply the uniform concentration result of [Zhou et al. \(2023b, Lemma 2\)](#) for the oracle influence difference function to obtain the following lemma.

Lemma A.2. *Suppose that Assumptions 2.1, 2.2, 2.3, and 4.2 hold. Then, for any stage $t \in \{1, 2, \dots, T\}$ and $\delta \in (0, 1)$, with probability at least $1 - 2\delta$, the following holds:*

$$\begin{aligned}
&\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} \left| \Delta_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) \right| \\
&\leq \left(54.4\sqrt{2}\kappa(\Pi_{t:T}) + 435.2 + \sqrt{2 \log \frac{1}{\delta}} \right) \sqrt{\frac{M_{t:T}^*}{n}} + o\left(\frac{1}{\sqrt{n}}\right),
\end{aligned} \tag{A.7}$$

where $M_{t:T}^* \equiv M \cdot \left(1 + 2\eta^{-T+t-1} + \sum_{s=1}^{T-t} 3\eta^{-s} \right) < \infty$.

Proof. See Appendix B.1. □

As for the latter (equation (A.6)), by extending the analytical strategy of Athey and Wager (2021) and Zhou et al. (2023b), which leverages orthogonal moments and cross-fitting, to our sequential multi-stage setting, we obtain the following lemma.

Lemma A.3. *Suppose that Assumptions 2.1, 2.2, 2.3, and 4.1 hold. Then, for any stage $t \in \{1, 2, \dots, T\}$, the following holds:*

$$\begin{aligned} & \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} |\widehat{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \widetilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T})| \\ & = O_p(n^{-\min\{1/2, \tau/2\}}). \end{aligned}$$

Proof. See Appendix B.3. □

Proof of Theorem 4.3. Combining the inequality (A.4) with Lemmas A.2 and A.3, we obtain

$$R_t^{\widehat{\pi}_{t:T}}(\pi_t^{*,B}) = O_p(\kappa(\Pi_{t:T}) \cdot n^{-1/2}) + O_p(n^{-\min\{1/2, \tau/2\}})$$

for all $t = 1, \dots, T$. This result proves Theorem 4.3 via the inequality (A.1) in Lemma A.1. □

B Preliminary Results and Proofs of Lemmas 3.1, 4.2, A.1, A.2, and A.3

B.1 Preliminary Results and Proofs of Lemmas 3.1, 4.2, and A.2

This section presents several preliminary results and proofs of Lemmas 3.1, 4.2, and A.2. We begin by providing the proof of Lemma 4.2.

Lemma B.1 below establishes a connection between the ϵ -Hamming covering numbers of classes for stage-specific policies and a class for sequences of policies, and will be used to prove Lemma 4.2.

Lemma B.1. *Given a class of DTRs $\Pi = \Pi_1 \times \cdots \times \Pi_T$, for any integers s and t such that $1 \leq s \leq t \leq T$, the following inequality holds:*

$$N_{d_h}((t-s+1)\epsilon, \Pi_{s:t}) \leq \prod_{\ell=s}^t N_{d_h}(\epsilon, \Pi_\ell).$$

Proof. Fix a set of history points $\{h_t^{(1)}, \dots, h_t^{(n)}\} \subseteq \mathcal{H}_h$. For any integer ℓ ($\leq t$), let $h_\ell^{(i)} (\subseteq h_t^{(i)})$ be the partial history up to stage ℓ . Let $K_\ell \equiv N_{d_h}(\epsilon, \Pi_\ell, \{h_\ell^{(1)}, \dots, h_\ell^{(n)}\})$. For each $\ell \in \{s, \dots, t\}$, we denote by $\tilde{\Pi}_\ell \equiv \left(\pi_\ell^{(1)}, \dots, \pi_\ell^{(K_\ell)}\right)$ the set of policies such that for any $\pi_\ell \in \Pi_\ell$, there exists $\pi_\ell^{(i)} \in \tilde{\Pi}_\ell$ satisfying $d_h(\pi_\ell, \pi_\ell^{(i)}) \leq \epsilon$. Such a set of policies exists from the definition of $N_{d_h}(\epsilon, \Pi_\ell, \{h_\ell^{(1)}, \dots, h_\ell^{(n)}\})$.

Fix $\pi_{s:t} \in \Pi_{s:t}$, and define $\tilde{\Pi}_{s:t} \equiv \tilde{\Pi}_s \times \cdots \times \tilde{\Pi}_t$. Let $\tilde{\pi}_{s:t} = (\tilde{\pi}_s, \dots, \tilde{\pi}_t) \in \tilde{\Pi}_{s:t}$ be such that for any $\ell \in \{s, \dots, t\}$, $d_h(\pi_\ell, \tilde{\pi}_\ell) \leq \epsilon$. Then

$$\begin{aligned} d_h(\pi_{s:t}, \tilde{\pi}_{s:t}) &= \frac{1}{n} \sum_{i=1}^n 1\{\pi_s(h_s^{(i)}) \neq \tilde{\pi}_s(h_s^{(i)}) \vee \cdots \vee \pi_t(h_t^{(i)}) \neq \tilde{\pi}_t(h_t^{(i)})\} \\ &\leq \sum_{\ell=s}^t \left(\frac{1}{n} \sum_{i=1}^n 1\{\pi_\ell(h_\ell^{(i)}) \neq \tilde{\pi}_\ell(h_\ell^{(i)})\} \right) \\ &= \sum_{\ell=s}^t d_h(\pi_\ell, \tilde{\pi}_\ell) \leq (t-s+1)\epsilon. \end{aligned}$$

Therefore, for any $\pi_{s:t} \in \Pi_{s:t}$, there exists $\tilde{\pi}_{s:t} \in \tilde{\Pi}_{s:t}$ such that $d_h(\pi_{s:t}, \tilde{\pi}_{s:t}) \leq (t-s+1)\epsilon$. Since $|\tilde{\Pi}_{s:t}| = \prod_{\ell=s}^t |\tilde{\Pi}_\ell| = \prod_{\ell=s}^t N_{d_h}(\epsilon, \Pi_\ell, \{h_\ell^{(1)}, \dots, h_\ell^{(n)}\})$, where $|\cdot|$ denotes the cardinality, we have

$$N_{d_h}((t-s+1)\epsilon, \Pi_{s:t}, \{h_t^{(1)}, \dots, h_t^{(n)}\}) \leq \prod_{\ell=s}^t N_{d_h}(\epsilon, \Pi_\ell, \{h_\ell^{(1)}, \dots, h_\ell^{(n)}\}).$$

As this holds for any n and any set of history points $\{h_t^{(1)}, \dots, h_t^{(n)}\}$, the result in the statement holds. \square

Using Lemma B.1, we present the proof of Lemma 4.2 below.

Proof of Lemma 4.2. Note that $\Pi = \Pi_{1:T}$. Applying Lemma B.1 to Π , we have

$N_H(\epsilon^2, \Pi) \leq \prod_{t=1}^T N_H(\epsilon^2/T, \Pi_t)$. Then

$$\begin{aligned}
\kappa(\Pi) &= \int_0^1 \sqrt{\log N_H(\epsilon^2, \Pi)} d\epsilon \leq \int_0^1 \sqrt{\sum_{t=1}^T \log N_H(\epsilon^2/T, \Pi_t)} d\epsilon \\
&\leq \sum_{t=1}^T \int_0^1 \sqrt{\log N_H(\epsilon^2/T, \Pi_t)} d\epsilon \\
&\leq T \int_0^1 \sqrt{\log C + D \left(\frac{\sqrt{T}}{\epsilon}\right)^{2\omega}} d\epsilon \\
&\leq T \int_0^1 \sqrt{\log C} d\epsilon + T \int_0^1 \sqrt{D \left(\frac{\sqrt{T}}{\epsilon}\right)^{2\omega}} d\epsilon \\
&= T\sqrt{\log C} + \sqrt{T^{(2+\omega)}}\sqrt{D} \int_0^1 \epsilon^{-\omega} d\epsilon = T\sqrt{\log C} + \frac{\sqrt{T^{(2+\omega)}}D}{1-\omega} \\
&< \infty,
\end{aligned}$$

where the third and last lines follow from Assumption 4.2. \square

We next give several preliminary results. We first define the conditional policy value of $\pi_{t:T}$ for stage t as, for any $h_t = (\underline{a}_{t-1}, \underline{s}_t) \in \mathcal{H}_t$,

$$V_t(\pi_{t:T}; h_t) \equiv E \left[\sum_{s=t}^T \sum_{\underline{a}'_{t:s} \in \underline{\mathcal{A}}_{t:s}} \left(Y_s(\underline{a}_{t-1}, \underline{a}'_{t:s}) \cdot \prod_{\ell=t}^s 1\{\pi_\ell(H_\ell(\underline{a}_{t-1}, \underline{a}'_{t:(\ell-1)})) = a'_\ell\} \right) \middle| H_t = h_t \right],$$

where $H_\ell(\underline{a}_{t-1}, \underline{a}'_{t:(\ell-1)}) = H_t(\underline{a}_{t-1})$ for $\ell = t$. Note that $E[V_t(\pi_{t:T}; H_t)] = V_t(\pi_{t:T})$. For ease of notations, for any stages s and t such that $s \geq t$ and any policy sequence $\pi_{t:s} \in \Pi_{t:s}$, we define

$$\tilde{Y}_{t:s}(\underline{a}_{t-1}, \pi_{t:s}) \equiv \sum_{\underline{a}'_{t:s} \in \underline{\mathcal{A}}_{t:s}} \left(Y_s(\underline{a}_{t-1}, \underline{a}'_{t:s}) \cdot \prod_{\ell=t}^s 1\{\pi_\ell(H_\ell(\underline{a}_{t-1}, \underline{a}'_{t:(\ell-1)})) = a'_\ell\} \right).$$

Let $\tilde{Y}_t(\underline{a}_{t-1}, \pi_t)$ denote $\tilde{Y}_{t:s}(\underline{a}_{t-1}, \pi_{t:s})$ when $t = s$. Note that, for any $\pi_{t:T} \in \Pi_{t:T}$ and $h_t = (\underline{a}_{t-1}, \underline{s}_t) \in \mathcal{H}_t$, $V(\pi_{t:T}; h_t)$ can be written as

$$E \left[\sum_{s=t}^T \tilde{Y}_{t:s}(\underline{a}_{t-1}, \pi_{t:s}) \middle| H_t = h_t \right] = V_t(\pi_{t:T}; h_t).$$

The following lemma will be used in the proofs of Lemmas 3.1, A.2, and A.3.

Lemma B.2. *Suppose that Assumption 2.1 holds. Then, for any stage t and DTR $\pi \in \Pi$, the following hold:*

$$(i) \quad Q_t^{\pi(t+1):T}(h_t, \pi_t) = V_t(\pi_{t:T}; h_t) \text{ for any } h_t \in \mathcal{H}_t;$$

$$(ii) \quad E \left[\tilde{V}_{i,t}(\pi_{t:T}) \right] = V_t(\pi_{t:T}) \text{ for any } i = 1, \dots, n.$$

Proof. We first prove (i), where we basically follow the proof of Tsiatis et al. (2019, equation (6.53)). For $t = T$, the following holds for any $h_T = (\underline{a}_{T-1}, \underline{s}_T) \in \mathcal{H}_T$:

$$\begin{aligned} Q_T(h_T, \pi_T) &= E[Y_T | A_T = \pi_T(h_T), H_T = h_T] \\ &= E[Y_T(\underline{A}_T) | A_T = \pi_T(h_T), H_T = h_T] \\ &= E[Y_T(\underline{a}_{T-1}, \pi_T(h_T)) | A_T = \pi_T(h_T), H_T = h_T] \\ &= \sum_{a_T \in \mathcal{A}_T} E[Y_T(\underline{a}_{T-1}, a_T) \cdot 1\{a_T = \pi_T(h_T)\} | A_T = a_T, H_T = h_T] \\ &= \sum_{a_T \in \mathcal{A}_T} E[Y_T(\underline{a}_{T-1}, a_T) \cdot 1\{a_T = \pi_T(h_T)\} | H_T = h_T] \\ &= E \left[\sum_{a_T \in \mathcal{A}_T} Y_T(\underline{a}_{T-1}, a_T) \cdot 1\{a_T = \pi_T(H_T(\underline{a}_{T-1}))\} \middle| H_T = h_T \right] \\ &= V_T(\pi_T; h_T), \end{aligned} \tag{B.1}$$

where the fifth equality follows from Assumption 2.1.

For any integer t such that $t < T$, the following holds for any $h_t = (\underline{a}_{t-1}, \underline{s}_t) \in \mathcal{H}_t$:

$$\begin{aligned} &E[Y_t + V_{t+1}(\pi_{(t+1):T}; H_{t+1}) | A_t = \pi_t(h_t), H_t = h_t] \\ &= E \left[Y_t(\underline{a}_t) + \sum_{s=t+1}^T \tilde{Y}_{(t+1):s}(\underline{a}_{t-1}, \pi_t(h_t), \pi_{(t+1):s}) \middle| A_t = \pi_t(h_t), H_t = h_t \right] \\ &= \sum_{a_t \in \mathcal{A}_t} \left(E \left[Y_t(\underline{a}_{t-1}, a_t) + \sum_{s=t+1}^T \tilde{Y}_{(t+1):s}(\underline{a}_{t-1}, a_t, \pi_{(t+1):s}) \middle| A_t = a_t, H_t = h_t \right] \cdot 1\{\pi_t(h_t) = a_t\} \right) \\ &= \sum_{a_t \in \mathcal{A}_t} \left(E \left[Y_t(\underline{a}_{t-1}, a_t) + \sum_{s=t+1}^T \tilde{Y}_{(t+1):s}(\underline{a}_{t-1}, a_t, \pi_{(t+1):s}) \middle| H_t = h_t \right] \cdot 1\{\pi_t(h_t) = a_t\} \right) \\ &= E \left[Y_t(\underline{a}_{t-1}, \pi_t(h_t)) + \sum_{s=t+1}^T \tilde{Y}_{(t+1):s}(\underline{a}_{t-1}, \pi_t(h_t), \pi_{(t+1):s}) \middle| H_t = h_t \right] \\ &= E \left[Y_t(\underline{a}_{t-1}, \pi_t(h_t)) + \sum_{s=t+1}^T \tilde{Y}_{(t+1):s}(\underline{A}_{t-1}, \pi_t(h_t), \pi_{(t+1):s}) \middle| H_t = h_t \right] \end{aligned}$$

$$\begin{aligned}
&= E \left[\sum_{s=t}^T \tilde{Y}_s(\underline{a}_{t-1}, \pi_{t:s}) \middle| H_t = h_t \right] \\
&= V_t(\pi_{t:T}; h_t),
\end{aligned} \tag{B.2}$$

where the first equality follows from the law of total expectations, and the third equality follows from Assumption 2.1.

When $t = T - 1$,

$$\begin{aligned}
Q_{T-1}^{\pi_T}(h_{T-1}, \pi_{T-1}) &= E[Y_{T-1} + Q_T(H_T, \pi_T) | A_{T-1} = \pi_{T-1}(h_{T-1}), H_{T-1} = h_{T-1}] \\
&= E[Y_{T-1} + V_T(\pi_T; H_T) | A_{T-1} = \pi_{T-1}(h_{T-1}), H_{T-1} = h_{T-1}] \\
&= V_{T-1}(\pi_{(T-1):T}; h_{T-1}),
\end{aligned}$$

where the second and third equalities follow from equations (B.1) and (B.2), respectively. Recursively applying the same argument from $t = T - 2$ to 1, we have $Q_t^{\pi^{(t+1):T}}(h_t, \pi_t) = V_t(\pi_{t:T}; h_t)$ for any t and $h_t \in \mathcal{H}_t$, which leads to the result (i).

We proceed to prove (ii). Given a fixed DTR $\pi = (\pi_1, \dots, \pi_T)$, to simplify notation, we define the independent copies of $\{\tilde{V}_{i,t}(\pi_{t:T}) : t = 1, \dots, T\}$ as follows:

$$\tilde{V}_T(\pi_T) \equiv \frac{Y_T - Q_T(H_T, A_T)}{e_T(H_T, A_T)} \cdot \mathbf{1}\{A_T = \pi_T(H_T)\} + Q_T(H_T, \pi_T(H_T)),$$

and, recursively for $t = T - 1, \dots, 1$,

$$\begin{aligned}
\tilde{V}_t(\pi_{t:T}) &\equiv \frac{Y_t + \tilde{V}_{t+1}(\pi_{(t+1):T}) - Q_t^{\pi^{(t+1):T}}(H_t, A_t)}{e_t(H_t, A_t)} \cdot \mathbf{1}\{A_t = \pi_t(H_t)\} \\
&\quad + Q_t^{\pi^{(t+1):T}}(H_t, \pi_t(H_t)).
\end{aligned}$$

Note that $E[\tilde{V}_T(\pi_T)] = E[\tilde{V}_{i,T}(\pi_T)]$ and $E[\tilde{V}_t(\pi_{t:T})] = E[\tilde{V}_{i,t}(\pi_{t:T})]$ for any i and t .

We first consider the case that $t = T$. Regarding the first component in $\tilde{V}_T(\pi_T)$, for any $h_T \in \mathcal{H}_T$,

$$\begin{aligned}
&E \left[Y_T \cdot \frac{\mathbf{1}\{A_T = \pi_T(H_T)\}}{e_T(H_T, A_T)} \middle| H_T = h_T \right] \\
&= E \left[E[Y_T | A_T = \pi_T(H_T), H_T] \cdot \frac{\mathbf{1}\{A_T = \pi_T(H_T)\}}{e_T(H_T, A_T)} \middle| H_T = h_T \right] \\
&= E \left[Q_T(H_T, \pi_T) \cdot \frac{\mathbf{1}\{A_T = \pi_T(H_T)\}}{e_T(H_T, A_T)} \middle| H_T = h_T \right],
\end{aligned}$$

where the last equality follows from the definition of $Q_T(H_T, \pi_T)$.

Therefore, we have

$$\begin{aligned}
& E \left[\tilde{V}_T(\pi_T) \middle| H_{i,T} = h_T \right] \\
&= E \left[(Q_T(H_T, A_T) - Q_T(H_T, A_T)) \cdot \frac{1\{A_T = \pi_T(H_T)\}}{e_T(H_T, A_T)} \middle| H_T = h_T \right] + Q_T(h_T, \pi_T) \\
&= Q_T(h_T, \pi_T) = V_T(\pi_T; h_T), \tag{B.3}
\end{aligned}$$

where the last equality follows from the result (i). We consequently have $E \left[\tilde{V}_T(\pi_T) \right] = E[V_T(\pi_T; H_T)] = V_T(\pi_T)$.

When $t = T - 1$, for any $h_{T-1} \in \mathcal{H}_{T-1}$,

$$\begin{aligned}
& E \left[\left(Y_{T-1} + \tilde{V}_T(\pi_T) \right) \cdot \frac{1\{A_{T-1} = \pi_{T-1}(H_{T-1})\}}{e_T(H_{T-1}, A_{T-1})} \middle| H_{T-1} = h_{T-1} \right] \\
&= E \left[\left(Y_{T-1} + E \left[\tilde{V}_T(\pi_T) \middle| H_T \right] \right) \cdot \frac{1\{A_{T-1} = \pi_{T-1}(H_{T-1})\}}{e_T(H_{T-1}, A_{T-1})} \middle| H_{T-1} = h_{T-1} \right] \\
&= E \left[\left(Y_{T-1} + V_T(\pi_T; H_T) \right) \cdot \frac{1\{A_{T-1} = \pi_{T-1}(H_{T-1})\}}{e_T(H_{T-1}, A_{T-1})} \middle| H_{T-1} = h_{T-1} \right] \\
&= E \left[\left(Y_{T-1} + Q_T(H_T, \pi_T) \right) \cdot \frac{1\{A_{T-1} = \pi_{T-1}(H_{T-1})\}}{e_T(H_{T-1}, A_{T-1})} \middle| H_{T-1} = h_{T-1} \right] \\
&= E \left[E \left[Y_{T-1} + Q_T(H_T, \pi_T) \middle| A_{T-1} = \pi_{T-1}(H_{T-1}), H_{T-1} \right] \right. \\
&\quad \left. \times \frac{1\{A_{T-1} = \pi_{T-1}(H_{T-1})\}}{e_T(H_{T-1}, A_{T-1})} \middle| H_{T-1} = h_{T-1} \right] \\
&= E \left[Q_{T-1}^{\pi_T}(H_{T-1}, A_{T-1}) \cdot \frac{1\{A_{T-1} = \pi_{T-1}(H_{T-1})\}}{e_T(H_{T-1}, A_{T-1})} \middle| H_{T-1} = h_{T-1} \right],
\end{aligned}$$

where the second equality follows from the result (B.3); the third equality follow from the result (i); the last equality follows from the definition of the Q-function $Q_{T-1}^{\pi_T}(\cdot, \cdot)$.

Therefore, for any $h_{T-1} \in \mathcal{H}_{T-1}$,

$$\begin{aligned}
& E \left[\tilde{V}_{T-1}(\pi_{(T-1):T}) \middle| H_{T-1} = h_{T-1} \right] \\
&= E \left[\left(Y_{T-1} + \tilde{V}_T(\pi_T) - Q_{T-1}^{\pi_T}(H_{T-1}, A_{T-1}) \right) \cdot \frac{1\{A_{T-1} = \pi_{T-1}(H_{T-1})\}}{e_T(H_{T-1}, A_{T-1})} \middle| H_{T-1} = h_{T-1} \right] \\
&\quad + Q_{T-1}^{\pi_T}(h_{T-1}, \pi_{T-1}) \\
&= E \left[\left(Q_{T-1}^{\pi_T}(H_{T-1}, A_{T-1}) - Q_{T-1}^{\pi_T}(H_{T-1}, A_{T-1}) \right) \cdot \frac{1\{A_{T-1} = \pi_{T-1}(H_{T-1})\}}{e_T(H_{T-1}, A_{T-1})} \middle| H_{T-1} = h_{T-1} \right] \\
&\quad + Q_{T-1}^{\pi_T}(h_{T-1}, \pi_{T-1})
\end{aligned}$$

$$= Q_{T-1}^{\pi_T}(h_{T-1}, \pi_{T-1}) = V_{T-1}(\pi_{(T-1):T}; h_{T-1}),$$

where the last equality follows from the result (i). Hence, we have $E \left[\tilde{V}_{T-1}(\pi_{(T-1):T}) \right] = E \left[V_{T-1}(\pi_{(T-1):T}; H_{T-1}) \right] = V_{T-1}(\pi_{(T-1):T})$.

Recursively applying the same argument from $t = T - 2$ to 1, we have $E \left[\tilde{V}_t(\pi_{t:T}) \right] = V_t(\pi_{t:T})$ for any t , which proves the result (iii). \square

We next provide the proof of Lemma 3.1, which extends the proof of Tsiatis et al. (2019, equation (7.21)) to the case under Assumption 3.1.

Proof of Lemma 3.1. Let $\pi \in \Pi$ be fixed. Under Assumptions 2.3 and 3.1, $\pi^{*,B}$ satisfies the condition in Assumption 3.1. For any $t = 1, \dots, T$, comparing the welfares of $(\pi_{1:(t-1)}, \pi_{t:T}^{*,B})$ and $(\pi_{1:t}, \pi_{(t+1):T}^{*,B})$ yields:

$$\begin{aligned} & W(\pi_{1:(t-1)}, \pi_{t:T}^{*,B}) - W(\pi_{1:t}, \pi_{(t+1):T}^{*,B}) \\ &= V_1(\pi_{1:(t-1)}, \pi_{t:T}^{*,B}) - V_1(\pi_{1:t}, \pi_{(t+1):T}^{*,B}) \\ &= E \left[Q_1^{\pi_{2:(t-1)}, \pi_{t:T}^{*,B}}(H_1, \pi_1) - Q_1^{\pi_{2:t}, \pi_{(t+1):T}^{*,B}}(H_1, \pi_1) \right] \\ &= E \left[E \left[Q_t^{\pi_{(t+1):T}^{*,B}}(H_t, \pi_t^{*,B}) - Q_t^{\pi_{(t+1):T}^{*,B}}(H_t, \pi_t) \middle| A_1 = \pi_1(H_1), \dots, A_{t-1} = \pi_{t-1}(H_{t-1}), H_{t-1} \right] \right] \\ &\geq 0, \end{aligned}$$

where we denote $W(\pi_{1:(0)}, \pi_{1:T}^{*,B}) = W(\pi^{*,B})$ and $W(\pi_{1:T}, \pi_{(T+1):T}^{*,B}) = W(\pi)$ for $t = 1$ and $t = T$, respectively; the second equality follows from Lemma B.2 (i); the third equality follows from the definitions of the Q-functions; the inequality follows from Assumption 3.1. We therefore obtain $W(\pi_{1:(t-1)}, \pi_{t:T}^{*,B}) \geq W(\pi_{1:t}, \pi_{(t+1):T}^{*,B})$ for any $t = 1, \dots, T$.

Applying this result, we have

$$W(\pi^{*,B}) - W(\pi) = \sum_{t=1}^T \left(W(\pi_{1:(t-1)}, \pi_{t:T}^{*,B}) - W(\pi_{1:t}, \pi_{(t+1):T}^{*,B}) \right) \geq 0.$$

Since this result holds for any $\pi \in \Pi$, we obtain the result in Lemma 3.1. \square

The following lemma, which follows from Lemma 2 in Zhou et al. (2023b) and its proof, plays important roles in the proofs of Theorem 4.3.

Lemma B.3. *Fix integers s and t such that $1 \leq s \leq t \leq T$. For any $\underline{a}_{s:t} \in \underline{\mathcal{A}}_{s:t}$, let $\{\Gamma_i^\dagger(\underline{a}_{s:t})\}_{i=1}^n$ be i.i.d. random variables with bounded supports. For any $\pi_{s:t} \in \Pi_{s:t}$,*

we define $\tilde{Q}(\pi_{s:t}) \equiv (1/n) \sum_{i=1}^n \Gamma_i^\dagger(\pi_{s:t})$, where $\Gamma_i^\dagger(\pi_{s:t}) \equiv \Gamma_i^\dagger((\pi_s(H_{i,s}), \dots, \pi_t(H_{i,t})))$ and $Q(\pi_{s:t}) \equiv E[\tilde{Q}(\pi_{s:t})]$. For any $\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}$, let $\tilde{\Delta}(\pi_{s:t}^a, \pi_{s:t}^b) \equiv \tilde{Q}(\pi_{s:t}^a) - \tilde{Q}(\pi_{s:t}^b)$ and $\Delta(\pi_{s:t}^a, \pi_{s:t}^b) \equiv Q(\pi_{s:t}^a) - Q(\pi_{s:t}^b)$. Then, when $\kappa(\Pi_{s:t}) < \infty$, the following holds: For any $\delta \in (0, 1)$, with probability at least $1 - 2\delta$,

$$\sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\Delta}(\pi_{s:t}^a, \pi_{s:t}^b) - \Delta(\pi_{s:t}^a, \pi_{s:t}^b) \right| \leq \left(54.4\sqrt{2}\kappa(\Pi_{s:t}) + 435.2 + \sqrt{2 \log \frac{1}{\delta}} \right) \sqrt{\frac{V_{s:t}^*}{n}} + o\left(\frac{1}{\sqrt{n}}\right),$$

where $V_{s:t}^* \equiv \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} E \left[\left(\Gamma_i^\dagger(\pi_{s:t}^a) - \Gamma_i^\dagger(\pi_{s:t}^b) \right)^2 \right] < \infty$.

Using Lemma B.3, we present the proof of Lemma A.2 as follows.

Proof of Lemma A.2. Given a fixed DTR $\pi = (\pi_1, \dots, \pi_T)$, we define the recursive expression for $t = T, \dots, 1$ as

$$\begin{aligned} \tilde{\Gamma}_{i,t}^{\pi^{(t+1):T}}(\underline{a}_{t:T}) &\equiv \frac{Y_{i,t} + \tilde{\Gamma}_{i,t+1}^{\pi^{(t+2):T}}(\underline{a}_{(t+1):T}) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, A_{i,t})}{e_t(H_{i,t}, A_{i,t})} \cdot 1\{A_{i,t} = a_t\} \\ &\quad + Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t), \end{aligned}$$

where we assume $\tilde{\Gamma}_{i,t}^{\pi^{(t+1):T}}(\underline{a}_{t:T}) = 0$ when $t = T + 1$. For any $\pi_{(t+1):T} \in \Pi_{(t+1):T}$ and $\pi'_{t:T} = (\pi'_t, \dots, \pi'_T) \in \Pi_{t:T}$, let $\tilde{\Gamma}_{i,t}^{\pi^{(t+1):T}}(\pi'_{t:T}) \equiv \tilde{\Gamma}_{i,t}^{\pi^{(t+1):T}}(\pi'_t(H_{i,t}), \dots, \pi'_T(H_{i,T}))$. Note that $\tilde{\Gamma}_{i,t}^{\pi^{(t+1):T}}(\pi_{t:T}) = \tilde{V}_{i,t}(\pi_{t:T})$.

Fix $\pi_{(t+1):T} \in \Pi_{(t+1):T}$. For any $\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}$, we define

$$\tilde{\Delta}^{\pi^{(t+1):T}}(\pi_{t:T}^a, \pi_{t:T}^b) \equiv \frac{1}{n} \sum_{i=1}^n \left(\tilde{\Gamma}_{i,t}^{\pi^{(t+1):T}}(\pi_{t:T}^a) - \tilde{\Gamma}_{i,t}^{\pi^{(t+1):T}}(\pi_{t:T}^b) \right)$$

and

$$\Delta^{\pi^{(t+1):T}}(\pi_{t:T}^a, \pi_{t:T}^b) \equiv E \left[\frac{1}{n} \sum_{i=1}^n \left(\tilde{\Gamma}_{i,t}^{\pi^{(t+1):T}}(\pi_{t:T}^a) - \tilde{\Gamma}_{i,t}^{\pi^{(t+1):T}}(\pi_{t:T}^b) \right) \right].$$

Note that

$$\tilde{\Delta}^{\pi^{(t+1):T}}(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) = \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}),$$

where $\tilde{\Delta}_t(\cdot; \cdot)$ is defined in (A.3). Noting that

$$\Delta^{\pi(t+1):T}(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) = E \left[\tilde{V}_{i,t}(\pi_t^a, \pi_{(t+1):T}) \right] - E \left[\tilde{V}_{i,t}(\pi_t^b, \pi_{(t+1):T}) \right],$$

Lemma B.2 leads to

$$\Delta^{\pi(t+1):T}(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) = \Delta_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}),$$

where $\Delta_t(\cdot; \cdot)$ is defined in (A.2).

Therefore, it follows for (A.5) that

$$\begin{aligned} & \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} |\Delta_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T})| \\ &= \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} |\tilde{\Delta}^{\pi(t+1):T}(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \Delta^{\pi(t+1):T}(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T})| \\ &\leq \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} |\tilde{\Delta}^{\pi(t+1):T}(\pi_{t:T}^a; \pi_{t:T}^b) - \Delta^{\pi(t+1):T}(\pi_{t:T}^a; \pi_{t:T}^b)|. \end{aligned} \quad (\text{B.4})$$

Fix $\pi_{(t+1):T} \in \Pi_{(t+1):T}$. Note that $\{\tilde{\Gamma}_{i,t}^{\pi(t+1):T}(\underline{a}_{t:T})\}_{i=1,\dots,n}$ are i.i.d. random variables with bounded supports under Assumptions 2.2 and 2.3 and that $\kappa(\Pi_{t:T})$ is finite from Lemma 4.2 and $\kappa(\Pi_{t:T}) \leq \kappa(\Pi)$. Therefore, fixing $\pi_{(t+1):T}$ and applying Lemma B.3 with $\Gamma_i^\dagger(\underline{a}_{t:T}) = \tilde{\Gamma}_{i,t}^{\pi(t+1):T}(\underline{a}_{t:T})$ leads to the following result: For any $\delta \in (0, 1)$, with probability at least $1 - 2\delta$,

$$\begin{aligned} & \sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} |\tilde{\Delta}^{\pi(t+1):T}(\pi_{t:T}^a; \pi_{t:T}^b) - \Delta^{\pi(t+1):T}(\pi_{t:T}^a; \pi_{t:T}^b)| \\ &\leq \left(54.4\sqrt{2}\kappa(\Pi_{t:T}) + 435.2 + \sqrt{2 \log \frac{1}{\delta}} \right) \sqrt{\frac{V_{t:T}^{\pi(t+1):T,*}}{n}} + o\left(\frac{1}{\sqrt{n}}\right), \end{aligned} \quad (\text{B.5})$$

with $V_{t:T}^{\pi(t+1):T,*} \equiv \sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} E \left[\left(\tilde{\Gamma}_{i,t}^{\pi(t+1):T}(\pi_{t:T}^a) - \tilde{\Gamma}_{i,t}^{\pi(t+1):T}(\pi_{t:T}^b) \right)^2 \right]$.

Under Assumptions 2.2 and 2.3, $V_{t:T}^{\pi(t+1):T,*} \leq M_{t:T}^* < \infty$ for any $\pi_{(t+1):T}$. Therefore, combining (B.4) and (B.5) leads to the result (A.7). \square

B.2 Proof of Lemma A.1

We provide the proof of Lemma A.1 in this section. The following lemma is a general version of Lemma A.1.

Lemma B.4. Fix $\pi = (\pi_1, \dots, \pi_T) \in \Pi$. Let $R_t^{\pi_{t:T}}(\tilde{\pi}_t) \equiv V_t(\tilde{\pi}_t, \pi_{(t+1):T}) - V_t(\pi_{t:T})$ for any $\tilde{\pi}_t \in \Pi_t$. Then, under Assumptions 2.1, 2.3, and 3.1, the regret of π is bounded from above as

$$R(\pi) \leq R_1^{\pi_{1:T}}(\pi_1^{*,B}) + \sum_{t=2}^T \frac{2^{t-2}}{\eta^{t-1}} R_t^{\pi_{t:T}}(\pi_t^{*,B}).$$

Proof. For any t , define $R_t(\pi_{t:T}) \equiv V_t(\pi_{t:T}^{*,B}) - V_t(\pi_{t:T})$, which is a partial regret of $\pi_{t:T}$ for stage t . Under Assumptions 2.3 and 3.1, $\pi^{*,B}$ satisfies the condition in Assumption 3.1. For any integers s and t such that $1 \leq t < s \leq T$,

$$\begin{aligned} & V_t(\pi_t^{*,B}, \dots, \pi_T^{*,B}) - V_t(\pi_t^{*,B}, \dots, \pi_{s-1}^{*,B}, \pi_s, \dots, \pi_T) \\ &= E \left[\frac{\prod_{\ell=t}^{s-1} 1\{A_\ell = \pi_\ell^{*,B}(H_\ell)\}}{\prod_{\ell=t}^{s-1} e_\ell(H_\ell, A_\ell)} \cdot \left(Q_s^{\pi_{(s+1):T}^{*,B}}(H_s, \pi_s^{*,B}) - Q_s^{\pi_{(s+1):T}}(H_s, \pi_s) \right) \right] \\ &\leq \frac{1}{\eta^{s-t}} E \left[\left(Q_s^{\pi_{(s+1):T}^{*,B}}(H_s, \pi_s^{*,B}) - Q_s^{\pi_{(s+1):T}}(H_s, \pi_s) \right) \right] \\ &= \frac{1}{\eta^{s-t}} \left(V_s(\pi_{s:T}^{*,B}) - V_s(\pi_{s:T}) \right) \\ &= \frac{1}{\eta^{s-t}} R_s(\pi_{t:T}), \end{aligned} \tag{B.6}$$

where the first equality follows from Lemma B.2 (i) and Assumption 2.1; the inequality follows from Assumptions 2.3 and 3.1; the second equality follows from Lemma B.2 (i).

For $t = T$ and $T - 1$, we have

$$\begin{aligned} R_T(\pi_T) &= V_T(\pi_T^{*,B}) - V_T(\pi_T) = R_T^{\pi_T}(\pi_T^{*,B}); \\ R_{T-1}(\pi_{(T-1):T}) &= \left[V_{T-1}(\pi_{T-1}^{*,B}, \pi_T^{*,B}) - V_{T-1}(\pi_{T-1}^{*,B}, \pi_T) \right] + \left[V_{T-1}(\pi_{T-1}^{*,B}, \pi_T) - V_{T-1}(\pi_{T-1}, \pi_T) \right] \\ &\leq \frac{1}{\eta} R_T^{\pi_T}(\pi_T^{*,B}) + R_{T-1}^{\pi_{(T-1):T}}(\pi_{T-1}^{*,B}) \\ &= \frac{1}{\eta} R_T(\pi_T) + R_{T-1}^{\pi_{(T-1):T}}(\pi_{T-1}^{*,B}), \end{aligned}$$

where the inequality follows from (B.6).

Generally, for $k = 2, \dots, T - 1$, it follows that

$$\begin{aligned} & R_{T-k}(\pi_{(T-k):T}) \\ &= V_{T-k}(\pi_{T-k}^{*,B}, \dots, \pi_T^{*,B}) - V_{T-k}(\pi_{T-k}, \dots, \pi_T) \end{aligned}$$

$$\begin{aligned}
&= \sum_{s=T-k}^T \left[V_{T-k} \left(\pi_{T-k}^{*,B}, \dots, \pi_s^{*,B}, \pi_{s+1}, \dots, \pi_T \right) - V_{T-k} \left(\pi_{T-k}^{*,B}, \dots, \pi_{s-1}^{*,B}, \pi_s, \dots, \pi_T \right) \right] \\
&= \sum_{s=T-k+1}^T \left[V_{T-k} \left(\pi_{T-k}^{*,B}, \dots, \pi_s^{*,B}, \pi_{s+1}, \dots, \pi_T \right) - V_{T-k} \left(\pi_{T-k}^{*,B}, \dots, \pi_{s-1}^{*,B}, \pi_s, \dots, \pi_T \right) \right] \\
&+ R_{T-k}^{\pi_{(T-k):T}}(\pi_{T-k}^{*,B}) \\
&\leq \sum_{s=T-k+1}^T \left[V_{T-k} \left(\pi_{T-k}^{*,B}, \dots, \pi_T^{*,B} \right) - V_{T-k} \left(\pi_{T-k}^{*,B}, \dots, \pi_{s-1}^{*,B}, \pi_s, \dots, \pi_T \right) \right] + R_{T-k}^{\pi_{(T-k):T}}(\pi_{T-k}^{*,B}) \\
&\leq \sum_{s=T-k+1}^T \frac{1}{\eta^{s-T+k}} R_s(\pi_{t:T}) + R_{T-k}^{\pi_{(T-k):T}}(\pi_{T-k}^{*,B}),
\end{aligned}$$

where the second equality follows from the telescoping sum; the third equality follows from the definition of $R_{T-k}^{\pi_{(T-k):T}}(\pi_{T-k}^{*,B})$; the first inequality follows from Lemma B.2 and Assumption 3.1; the last line follows from (B.6).

Then, recursively, the following hold:

$$\begin{aligned}
R_{T-1}(\pi_{(T-1):T}) &\leq \frac{1}{\eta} R_T(\pi_T) + R_{T-1}^{\pi_{(T-1):T}}(\pi_{T-1}^{*,B}) \\
&= \frac{1}{\eta} R_T^{\pi_T}(\pi_T^{*,B}) + R_{T-1}^{\pi_{(T-1):T}}(\pi_{T-1}^{*,B}), \\
R_{T-2}(\pi_{(T-2):T}) &\leq \frac{1}{\eta} R_{T-1}(\pi_{(T-1):T}) + \frac{1}{\eta^2} R_T(\pi_T) + R_{T-2}^{\pi_{(T-2):T}}(\pi_{T-2}^{*,B}) \\
&\leq \frac{2}{\eta^2} R_T^{\pi_T}(\pi_T^{*,B}) + \frac{1}{\eta} R_{T-1}^{\pi_{(T-1):T}}(\pi_{T-1}^{*,B}) + R_{T-2}^{\pi_{(T-2):T}}(\pi_{T-2}^{*,B}), \\
&\vdots \\
R_{T-k}(\pi_{(T-k):T}) &\leq \sum_{s=1}^k \frac{2^{k-s}}{\eta^{k-s+1}} R_{T-s+1}^{\pi_{(T-s+1):T}}(\pi_{T-s+1}^{*,B}) + R_{T-k}^{\pi_{(T-k):T}}(\pi_{T-k}^{*,B}).
\end{aligned}$$

Therefore, setting $k = T - 1$ and noting that $R_1(\pi_{1:T}) = R(\pi)$, we obtain

$$\begin{aligned}
R(\pi) &\leq \sum_{s=1}^{T-1} \frac{2^{T-1-s}}{\eta^{T-s}} R_{T-s+1}^{\pi_{(T-s+1):T}}(\pi_{T-s+1}^{*,B}) + R_1^{\pi_{1:T}}(\pi_1^{*,B}) \\
&= R_1^{\pi_{1:T}}(\pi_1^{*,B}) + \sum_{s=1}^{T-1} \frac{2^{s-1}}{\eta^s} R_{s+1}^{\pi_{(s+1):T}}(\pi_{s+1}^{*,B}).
\end{aligned}$$

Setting $t = s + 1$ in the above equation leads to the result. \square

The proof of Lemma A.1 is given below.

Proof of Lemma A.1. Lemma A.1 follows immediately from Lemma B.4 with setting

$\pi = \hat{\pi}$. □

B.3 Proof of Lemma A.3

This section presents the proof of Lemma A.3. We here consider extending the analytical strategy of Athey and Wager (2021) and Zhou et al. (2023b), which leverages orthogonal moments and cross-fitting, to the sequential multi-stage setting.

We begin by introducing several notational definitions. For any $\underline{a}_{s:t} \in \underline{\mathcal{A}}_{s:t}$ and $\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}$, let

$$G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{\underline{a}_{s:t}} \equiv \prod_{\ell=s}^t 1\{\pi_{\ell}^a(H_{i,\ell}) = a_{\ell}\} - \prod_{\ell=s}^t 1\{\pi_{\ell}^b(H_{i,\ell}) = a_{\ell}\}.$$

Given a fixed DTR $\pi = (\pi_1, \dots, \pi_T)$, with some abuse of notation, we define

$$\begin{aligned} \tilde{V}_{i,T}(a_T) &\equiv \frac{Y_{i,T} - Q_T(H_{i,T}, a_T)}{e_T(H_{i,T}, a_T)} \cdot 1\{A_{i,T} = a_T\} + Q_T(H_{i,T}, a_T), \\ \hat{V}_{i,T}(a_T) &\equiv \frac{Y_{i,T} - \hat{Q}_T^{-k(i)}(H_{i,T}, a_T)}{\hat{e}_T^{-k(i)}(H_{i,T}, a_T)} \cdot 1\{A_{i,T} = a_T\} + \hat{Q}_T^{-k(i)}(H_{i,T}, a_T), \end{aligned}$$

and, recursively, for $t = T - 1, \dots, 1$,

$$\begin{aligned} \tilde{V}_{i,t}^{\pi^{(t+1):T}}(a_t) &\equiv \frac{Y_{i,t} + \tilde{V}_{i,t+1}^{\pi^{(t+2):T}}(\pi_{t+1}(H_{i,t+1})) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t)}{e_t(H_{i,t}, a_t)} \cdot 1\{A_{i,t} = a_t\} \\ &\quad + Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t), \\ \hat{V}_{i,t}^{\pi^{(t+1):T}}(a_t) &\equiv \frac{Y_{i,t} + \hat{V}_{i,t+1}^{\pi^{(t+2):T}}(\pi_{t+1}(H_{i,t+1})) - \hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t)}{\hat{e}_t^{-k(i)}(H_{i,t}, a_t)} \cdot 1\{A_{i,t} = a_t\} \\ &\quad + \hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t), \end{aligned}$$

where, for $t = T$, we denote $\tilde{V}_{i,T}^{\pi^{(T+1):T}}(a_t) = \tilde{V}_{i,T}(a_T)$ and $\hat{V}_{i,T}^{\pi^{(T+1):T}}(a_t) = \hat{V}_{i,T}(a_T)$. Note that the inside of (A.6) is expressed as

$$\begin{aligned} &\hat{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) \\ &= \sum_{a_t \in \mathcal{A}_t} G_{i, \pi_t^a, \pi_t^b}^{a_t} \left(\hat{V}_{i,t}^{\pi^{(t+1):T}}(a_t) - \tilde{V}_{i,t}^{\pi^{(t+1):T}}(a_t) \right). \end{aligned} \tag{B.7}$$

For integers s and t such that $1 \leq s \leq t \leq T$ and $\underline{a}_{s:t} \in \underline{\mathcal{A}}_{s:t}$, let

$$\begin{aligned} & \tilde{\mathcal{S}}_{s:t}^{\underline{a}_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \\ & \equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{\underline{a}_{s:t}} \cdot \left(\frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \right) \cdot \left(\hat{V}_{i,t}^{\pi_{(t+1):T}}(a_t) - \tilde{V}_{i,t}^{\pi_{(t+1):T}}(a_t) \right), \end{aligned}$$

where we assume that $(\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}) / (\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)) = 1$ when $s = t$. We denote $\tilde{\mathcal{S}}_t^{a_t}(\pi_t^a, \pi_t^b, \pi_{(t+1):T}) = \tilde{\mathcal{S}}_{t:t}^{a_t}(\pi_{t:t}^a, \pi_{t:t}^b, \pi_{(t+1):T})$ for the case $s = t$. For $t = T$, we also denote $\tilde{\mathcal{S}}_{s:T}^{\underline{a}_{s:T}}(\pi_{s:T}^a, \pi_{s:T}^b, \pi_{(t+1):T}) = \tilde{\mathcal{S}}_{s:T}^{\underline{a}_{s:T}}(\pi_{s:T}^a, \pi_{s:T}^b)$ where

$$\tilde{\mathcal{S}}_{s:T}^{\underline{a}_{s:T}}(\pi_{s:T}^a, \pi_{s:T}^b) \equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:T}^a, \pi_{s:T}^b}^{\underline{a}_{s:T}} \cdot \left(\frac{\prod_{\ell=s}^{T-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{T-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \right) \cdot \left(\hat{V}_{i,T}(a_T) - \tilde{V}_{i,T}(a_T) \right).$$

Note that

$$\sum_{a_t \in \mathcal{A}_t} \tilde{\mathcal{S}}_t^{a_t}(\pi_t^a, \pi_t^b, \pi_{(t+1):T}) = \hat{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T})$$

holds from equation (B.7). Hence, regarding (A.6), it follows that

$$\begin{aligned} & \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} |\hat{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T}) - \tilde{\Delta}_t(\pi_t^a, \pi_{(t+1):T}; \pi_t^b, \pi_{(t+1):T})| \\ & \leq \sum_{a_t \in \mathcal{A}_t} \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} |\tilde{\mathcal{S}}_t^{a_t}(\pi_t^a, \pi_t^b, \pi_{(t+1):T})|. \end{aligned} \quad (\text{B.8})$$

This result enables us to evaluate (A.6) through evaluating

$$\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} |\tilde{\mathcal{S}}_t^{a_t}(\pi_t^a, \pi_t^b, \pi_{(t+1):T})|$$

for each $a_t \in \mathcal{A}_t$.

Lemma B.5. *Suppose that Assumptions 2.1, 2.2, 2.3, 4.1, and 4.2 hold. Then for any integers s and t such that $1 \leq s \leq t < T$,*

$$\begin{aligned} & \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t}^{\underline{a}_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| \\ & \leq O_p(n^{-\min\{1/2, \tau/2\}}) \\ & + \sum_{a_{t+1} \in \mathcal{A}_{t+1}} \sup_{\pi_{(t+2):T} \in \Pi_{(t+2):T}} \sup_{\pi_{s:(t+1)}^a, \pi_{s:(t+1)}^b \in \Pi_{s:(t+1)}} \left| \tilde{\mathcal{S}}_{s:(t+1)}^{\underline{a}_{s:(t+1)}}(\pi_{s:(t+1)}^a, \pi_{s:(t+1)}^b, \pi_{(t+2):T}) \right|. \end{aligned} \quad (\text{B.9})$$

Moreover, for any integer t such that $1 \leq t \leq T$,

$$\sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} \left| \tilde{S}_{t:T}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b) \right| = O_p(n^{-\min\{1/2, \tau/2\}}). \quad (\text{B.10})$$

Proof. We first consider the case that $t < T$. For any integers s and t such that $1 \leq s \leq t < T$, we define

$$\begin{aligned} \tilde{S}_{s:t,(A)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) &\equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \\ &\quad \times \left(\hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,t}, a_t) - Q_t^{\pi_{(t+1):T}}(H_{i,t}, a_t) \right) \\ &\quad \times \left(1 - \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right); \\ \tilde{S}_{s:t,(B)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) &\equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \\ &\quad \times \left(\tilde{V}_{i,t+1}^{\pi_{(t+2):T}}(\pi_{t+1}(H_{i,t+1})) - Q_t^{\pi_{(t+1):T}}(H_{i,t}, a_t) \right) \\ &\quad \times \left(\frac{1\{A_{i,t} = a_t\}}{\hat{e}_t^{-k(i)}(H_{i,t}, a_t)} - \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right); \\ \tilde{S}_{s:t,(C)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) &\equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \\ &\quad \times \left(Q_t^{\pi_{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right) \\ &\quad \times \left(\frac{1\{A_{i,t} = a_t\}}{\hat{e}_t^{-k(i)}(H_{i,t}, a_t)} - \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right). \end{aligned}$$

Using these definitions, we can decompose $\tilde{S}_{s:t}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T})$ as follows:

$$\begin{aligned} &\tilde{S}_{s:t}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \\ &= \tilde{S}_{s:t,(A)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) + \tilde{S}_{s:t,(B)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) + \tilde{S}_{s:t,(C)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \\ &\quad + \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \left(\frac{\prod_{\ell=s}^t 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^t \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \right) \left(\hat{V}_{i,t+1}^{\pi_{(t+2):T}}(\pi_{t+1}(H_{i,t+1})) - \tilde{V}_{i,t+1}^{\pi_{(t+2):T}}(\pi_{t+1}(H_{i,t+1})) \right). \end{aligned}$$

Regarding the forth term in the above equation, it follows that

$$\begin{aligned} &\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \left(\frac{\prod_{\ell=s}^t 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^t \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \right) \right. \\ &\quad \left. \times \left(\hat{V}_{i,t+1}^{\pi_{(t+2):T}}(\pi_{t+1}(H_{i,t+1})) - \tilde{V}_{i,t+1}^{\pi_{(t+2):T}}(\pi_{t+1}(H_{i,t+1})) \right) \right| \end{aligned}$$

$$\leq \sum_{a_{t+1} \in \mathcal{A}_{t+1}} \sup_{\pi^{(t+2):T} \in \Pi^{(t+2):T}} \sup_{\pi_{s:(t+1)}^a, \pi_{s:(t+1)}^b \in \Pi_{s:(t+1)}} \left| \tilde{\mathcal{S}}_{s:(t+1)}^{a_s:(t+1)}(\pi_{s:(t+1)}^a, \pi_{s:(t+1)}^b, \pi^{(t+2):T}) \right|.$$

Hence, we have

$$\begin{aligned} & \sup_{\pi^{(t+1):T} \in \Pi^{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) \right| \\ & \leq \sup_{\pi^{(t+1):T} \in \Pi^{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t,(A)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) \right| \\ & + \sup_{\pi^{(t+1):T} \in \Pi^{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t,(B)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) \right| \\ & + \sup_{\pi^{(t+1):T} \in \Pi^{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t,(C)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) \right| \\ & + \sum_{a_{t+1} \in \mathcal{A}_{t+1}} \sup_{\pi^{(t+2):T} \in \Pi^{(t+2):T}} \sup_{\pi_{s:(t+1)}^a, \pi_{s:(t+1)}^b \in \Pi_{s:(t+1)}} \left| \tilde{\mathcal{S}}_{s:(t+1)}^{a_s:(t+1)}(\pi_{s:(t+1)}^a, \pi_{s:(t+1)}^b, \pi^{(t+2):T}) \right|. \quad (\text{B.11}) \end{aligned}$$

In what follows, we will prove that:

$$\sup_{\pi^{(t+1):T} \in \Pi^{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t,(A)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) \right| = O_P(n^{-1/2}); \quad (\text{B.12})$$

$$\sup_{\pi^{(t+1):T} \in \Pi^{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t,(B)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) \right| = O_P(n^{-1/2}); \quad (\text{B.13})$$

$$\sup_{\pi^{(t+1):T} \in \Pi^{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t,(C)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) \right| = O_P(n^{-\min\{1/2, \tau/2\}}). \quad (\text{B.14})$$

Then we can obtain the result (B.9) from equation (B.11).

Throughout the proof, without loss of generality, we assume that $n > n_0$, where n_0 is defined in Assumption 4.1. We begin by examining $\tilde{\mathcal{S}}_{s:t,(A)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T})$, which we decompose as follows:

$$\tilde{\mathcal{S}}_{s:t,(A)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) = \tilde{\mathcal{S}}_{s:t,(A1)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) + \tilde{\mathcal{S}}_{s:t,(A2)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}),$$

where

$$\begin{aligned} \tilde{\mathcal{S}}_{s:t,(A1)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) & \equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_s:t} \cdot \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \\ & \times \left(\hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right) \left(1 - \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right); \\ \tilde{\mathcal{S}}_{s:t,(A2)}^{a_s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) & \equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_s:t} \cdot \left(\frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right) \end{aligned}$$

$$\times \left(\widehat{Q}_t^{\pi^{(t+1):T}, -k^{(i)}}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right) \left(1 - \frac{1 \{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right).$$

For each fold index k , we define

$$\begin{aligned} \widetilde{S}_{s:t, (A1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) &\equiv \frac{1}{n} \sum_{i \in I_k} G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \cdot \frac{\prod_{\ell=s}^{t-1} 1 \{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \\ &\times \left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right) \left(1 - \frac{1 \{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right). \end{aligned}$$

Fix $k \in \{1, \dots, K\}$. We now consider $\widetilde{S}_{s:t, (A1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T})$. Since $\widehat{Q}_t^{\pi^{(t+1):T}, -k}(\cdot, a_t)$ is computed using the data in the rest $K-1$ folds, when the data $\mathcal{S}_{-k} \equiv \{Z_i : i \notin I_k\}$ in the rest $K-1$ folds is conditioned, $\widehat{Q}_t^{\pi^{(t+1):T}, -k}(\cdot, a_{t:T})$ is fixed; hence, $\widetilde{S}_{s:t, (A1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T})$ is a sum of i.i.d. bounded random variables under Assumptions 2.2, 2.3, and 4.1 (ii).

It follows that

$$\begin{aligned} &E \left[G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \cdot \frac{\prod_{\ell=s}^{t-1} 1 \{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \cdot \left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right) \right. \\ &\quad \left. \times \left(1 - \frac{1 \{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right) \middle| \mathcal{S}_{-k} \right] \\ &= E \left[G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \cdot \frac{\prod_{\ell=s}^{t-1} 1 \{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \cdot \left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right) \right. \\ &\quad \left. \times E \left[\left(1 - \frac{1 \{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right) \middle| H_{i,t} \right] \middle| \mathcal{S}_{-k} \right] \\ &= E \left[G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \cdot \frac{\prod_{\ell=s}^{t-1} 1 \{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \cdot \left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right) \right. \\ &\quad \left. \times \left(1 - \frac{e_t(H_{i,t}, a_t)}{e_t(H_{i,t}, a_t)} \right) \middle| \mathcal{S}_{-k} \right] \\ &= 0. \end{aligned}$$

Hence, fixing $\pi_{(t+1):T}$, $\sup_{\pi_{s:T}^a, \pi_{s:T}^b \in \Pi_{s:t}} \left| \widetilde{S}_{s:t, (A1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right|$ can be written as

$$\begin{aligned} &\sup_{\pi_{s:T}^a, \pi_{s:T}^b \in \Pi_{s:t}} \left| \widetilde{S}_{s:t, (A1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| \\ &= \frac{1}{K} \sup_{\pi_{s:T}^a, \pi_{s:T}^b \in \Pi_{s:t}} \left| \frac{1}{n/K} \sum_{i \in I_k} G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \cdot \frac{\prod_{\ell=s}^{t-1} 1 \{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right. \\ &\quad \left. \times \left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right) \left(1 - \frac{1 \{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right) \right| \end{aligned}$$

$$\begin{aligned}
& - E \left[\frac{1}{n/K} \sum_{i \in I_k} G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \cdot \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right. \\
& \quad \left. \times \left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right) \left(1 - \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right) \middle| \mathcal{S}_{-k} \right].
\end{aligned}$$

By applying Lemma B.3 while fixing \mathcal{S}_{-k} and setting $i \in I_k$ and

$$\Gamma_i^\dagger(a_{s:t}) = \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right) \left(1 - \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right),$$

the following holds: $\forall \delta > 0$, with probability at least $1 - 2\delta$,

$$\begin{aligned}
& \sup_{\pi_{s:T}^a, \pi_{s:T}^b \in \Pi_{s:t}} \left| \widetilde{S}_{s:t, (A1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) \right| \\
& \leq o(n^{-1/2}) + \left(54.4\kappa(\Pi_{s:t}) + 435.2 + \sqrt{2 \log(1/\delta)} \right) \\
& \quad \times \left[\sup_{\pi_{s:T} \in \Pi_{s:T}} E \left[\left(G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \right)^2 \cdot \left(\frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right)^2 \right. \right. \\
& \quad \left. \left. \times \left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right)^2 \left(1 - \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right)^2 \middle| \mathcal{S}_{-k} \right] \middle/ \left(\frac{n}{K} \right) \right]^{1/2} \\
& \leq o(n^{-1/2}) + \sqrt{K} \cdot \left(54.4\kappa(\Pi_{s:t}) + 435.2 + \sqrt{2 \log(1/\delta)} \right) \cdot \left(\frac{1}{\eta} \right)^{t-s+1} \\
& \quad \times \sqrt{\frac{E \left[\left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right)^2 \middle| \mathcal{S}_{-k} \right]}{n}},
\end{aligned}$$

where the last inequality follows from $\left(G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \right)^2 \leq 1$ a.s. and Assumption 2.3 (overlap condition). Note that $\kappa(\Pi_{s:t})$ is finite by Lemma 4.2 and the inequality $\kappa(\Pi_{s:t}) \leq \kappa(\Pi)$.

From Assumptions 2.2 and 4.1 (ii), we have

$$\sup_{\pi^{(t+1):T} \in \Pi^{(t+1):T}} E \left[\left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right)^2 \right] < \infty.$$

Hence, Markov's inequality yields

$$\sup_{\pi^{(t+1):T} \in \Pi^{(t+1):T}} E \left[\left(\widehat{Q}_t^{\pi^{(t+1):T}, -k}(H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) \right)^2 \middle| \mathcal{S}_{-k} \right] = O_p(1).$$

Combining these results, we obtain

$$\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:T}^a, \pi_{s:T}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (A1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| = O_p \left(\frac{1}{\sqrt{n}} \right). \quad (\text{B.15})$$

Consequently,

$$\begin{aligned} & \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:T}^a, \pi_{s:T}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (A1)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| \\ & \leq \sum_{k=1}^K \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:T}^a, \pi_{s:T}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (A1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| = O_p \left(\frac{1}{\sqrt{n}} \right). \end{aligned}$$

We next consider $\tilde{S}_{s:t, (B)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T})$ (we will consider $\tilde{S}_{s:t, (A2)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T})$ later). We begin by decomposing $\tilde{S}_{s:t, (B)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T})$ as follows:

$$\tilde{S}_{s:t, (B)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) = \sum_{k=1}^K \left(\tilde{S}_{s:t, (B1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) + \tilde{S}_{s:t, (B2)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right),$$

where

$$\begin{aligned} \tilde{S}_{s:t, (B1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) & \equiv \frac{1}{n} \sum_{i \in I_k} G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \left(\tilde{V}_{i, t+1}^{\pi_{(t+2):T}}(\pi_{t+1}(H_{i, t+1})) - Q_t^{\pi_{(t+1):T}}(H_{i, t}, a_t) \right) \\ & \quad \times \left(\frac{\prod_{\ell=s}^{t-1} 1\{A_{i, \ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i, \ell}, a_\ell)} - \frac{\prod_{\ell=s}^{t-1} 1\{A_{i, \ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k}(H_{i, \ell}, a_\ell)} \right) \frac{1\{A_{i, t} = a_t\}}{e_t(H_{i, t}, a_t)}, \\ \tilde{S}_{s:t, (B2)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) & \equiv \frac{1}{n} \sum_{i \in I_k} G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \left(\tilde{V}_{i, t+1}^{\pi_{(t+2):T}}(\pi_{t+1}(H_{i, t+1})) - Q_t^{\pi_{(t+1):T}}(H_{i, t}, a_t) \right) \\ & \quad \times \left(\frac{\prod_{\ell=s}^t 1\{A_{i, \ell} = a_\ell\}}{\prod_{\ell=s}^t \hat{e}_\ell^{-k(i)}(H_{i, \ell}, a_\ell)} - \frac{\prod_{\ell=s}^t 1\{A_{i, \ell} = a_\ell\}}{\prod_{\ell=s}^t e_\ell(H_{i, \ell}, a_\ell)} \right). \end{aligned}$$

Fix $k \in \{1, \dots, K\}$. As for $\tilde{S}_{s:t, (B1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T})$, taking the conditional expectation given the date \mathcal{S}_{-k} in the rest $k-1$ folds leads to

$$\begin{aligned} & E \left[\tilde{S}_{s:t, (B1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \middle| \mathcal{S}_{-k} \right] \\ & = E \left[\frac{1}{n/K} \sum_{i \in I_k} G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \left(\tilde{V}_{i, t+1}^{\pi_{(t+2):T}}(\pi_{t+1}(H_{i, t+1})) - Q_t^{\pi_{(t+1):T}}(H_{i, t}, a_t) \right) \right. \\ & \quad \left. \times \left(\frac{\prod_{\ell=s}^{t-1} 1\{A_{i, \ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i, \ell}, a_\ell)} - \frac{\prod_{\ell=s}^{t-1} 1\{A_{i, \ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k}(H_{i, \ell}, a_\ell)} \right) \frac{1\{A_{i, t} = a_t\}}{e_t(H_{i, t}, a_t)} \middle| \mathcal{S}_{-k} \right] \end{aligned}$$

$$\begin{aligned}
&= E \left[\frac{1}{n/K} \sum_{i \in I_k} G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} E \left[\tilde{V}_{i,t+1}^{\pi^{(t+2):T}} (\pi_{t+1}(H_{i,t+1})) - Q_t^{\pi^{(t+1):T}} (H_{i,t}, a_t) \middle| H_{i,t}, A_{i,t} = a_t \right] \right. \\
&\quad \times \left. \left(\frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} - \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k}(H_{i,\ell}, a_\ell)} \right) \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \middle| \mathcal{S}_{-k} \right] \\
&= E \left[\frac{1}{n/K} \sum_{i \in I_k} G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} (Q_t^{\pi^{(t+1):T}} (H_{i,t}, a_t) - Q_t^{\pi^{(t+1):T}} (H_{i,t}, a_t)) \right. \\
&\quad \times \left. \left(\frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} - \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k}(H_{i,\ell}, a_\ell)} \right) \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \middle| \mathcal{S}_{-k} \right] \\
&= 0,
\end{aligned}$$

where the third equality follows from Lemma B.2 (i). Note that conditional on \mathcal{S}_{-k} , $\tilde{S}_{s:t, (B1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T})$ is a sum of i.i.d. bounded random variables under Assumptions 2.2, 2.3, and 4.1 (ii), and its condition mean is zero. Hence, fixing $\pi_{(t+1):T}$ and conditioning on \mathcal{S}_{-k} , we can apply Lemma B.3 with setting $i \in I_k$ and

$$\begin{aligned}
\Gamma_i^\dagger(a_{s:t}) &= \left(\tilde{V}_{i,t+1}^{\pi^{(t+2):T}} (\pi_{t+1}(H_{i,t+1})) - Q_t^{\pi^{(t+1):T}} (H_{i,t}, a_t) \right) \\
&\quad \times \left(\frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} - \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k}(H_{i,\ell}, a_\ell)} \right) \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)}
\end{aligned}$$

to obtain the following: $\forall \delta > 0$, with probability at least $1 - 2\delta$,

$$\begin{aligned}
&\sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (B1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| \\
&\leq o(n^{-1/2}) + \left(54.4\kappa(\Pi_{s:t}) + 435.2 + \sqrt{2 \log(1/\delta)} \right) \\
&\quad \times \left[\sup_{\pi_{s:T} \in \Pi_{s:T}} E \left[\left(G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \right)^2 \cdot \left(\tilde{V}_{i,t+1}^{\pi^{(t+2):T}} (\pi_{t+1}(H_{i,t+1})) - Q_t^{\pi^{(t+1):T}} (H_{i,t}, a_t) \right)^2 \right. \right. \\
&\quad \left. \left. \left(\frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} - \frac{\prod_{\ell=s}^{t-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k}(H_{i,\ell}, a_\ell)} \right)^2 \left(\frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right)^2 \middle| \mathcal{S}_{-k} \right] \middle/ \left(\frac{n}{K} \right) \right]^{1/2} \\
&\leq o(n^{-1/2}) + \sqrt{K} \cdot \left(54.4\kappa(\pi_{t:T}) + 435.2 + \sqrt{2 \log(1/\delta)} \right) \cdot \left(\sum_{j=0}^{T-t} \frac{3M}{\eta^j} \right) \\
&\quad \times \sqrt{\frac{E \left[\left(\frac{1}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k}(H_{i,\ell}, a_\ell)} \right)^2 \middle| \mathcal{S}_{-k} \right]}{n}},
\end{aligned}$$

where the last inequality follows from $\left(G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \right)^2 \leq 1$ a.s. and Assumptions 2.2 and 2.3.

From Assumptions 2.3 and 4.1 (ii), we have

$$E \left[\left(\frac{1}{\prod_{\ell=s}^{t-1} e_{\ell}(H_{i,\ell}, a_{\ell})} - \frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_{\ell}^{-k}(H_{i,\ell}, a_{\ell})} \right)^2 \right] < \infty.$$

Hence, Markov's inequality leads to

$$E \left[\left(\frac{1}{\prod_{\ell=s}^{t-1} e_{\ell}(H_{i,\ell}, a_{\ell})} - \frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_{\ell}^{-k}(H_{i,\ell}, a_{\ell})} \right)^2 \middle| \mathcal{S}_{-k} \right] = O_p(1).$$

Note also that $\kappa(\Pi_{s:t}) < \infty$ by Lemma 4.2. Combining these results, we have

$$\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t, (B1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| = O_p \left(\frac{1}{\sqrt{n}} \right). \quad (\text{B.16})$$

By applying the same argument to derive (B.16), we also obtain

$$\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t, (B2)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| = O_p \left(\frac{1}{\sqrt{n}} \right).$$

Consequently,

$$\begin{aligned} & \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t, (B)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| \\ & \leq \sum_{k=1}^K \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t, (B1)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| \\ & + \sum_{k=1}^K \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t, (B2)}^{a_{s:t}, k}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| \\ & = O_p \left(\frac{1}{\sqrt{n}} \right), \end{aligned}$$

which proves equation (B.13).

We next consider to bound $\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t, (C)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right|$ from above. It follows that

$$\begin{aligned} & \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{\mathcal{S}}_{s:t, (C)}^{a_{s:t}}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| \\ & = \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{s:t}^a, \pi_{s:t}^b}^{a_{s:t}} \frac{\prod_{\ell=s}^{t-1} \mathbf{1}\{A_{i,\ell} = a_{\ell}\}}{\prod_{\ell=s}^{t-1} \hat{e}_{\ell}^{-k(i)}(H_{i,\ell}, a_{\ell})} \left(Q_t^{\pi_{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right) \right| \end{aligned}$$

$$\begin{aligned}
& \times \left(\frac{1 \{A_{i,t} = a_t\}}{\hat{e}_t^{-k(i)}(H_{i,t}, a_t)} - \frac{1 \{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right) \Big| \\
& \leq \frac{1}{n} \sum_{i=1}^n \left| Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right| \\
& \times \left| \frac{1}{\prod_{\ell=s}^t \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{e_t(H_{i,t}, a_t) \cdot \prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \right| \\
& \leq \frac{1}{n} \sum_{i=1}^n \left| Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right| \cdot \left| \frac{1}{\prod_{\ell=s}^t \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^t e_\ell(H_{i,\ell}, a_\ell)} \right| \\
& + \frac{1}{n} \sum_{i=1}^n \left| Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right| \\
& \times \left| \frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right| \left(\frac{1}{e_t(H_{i,t}, a_t)} \right) \\
& \leq \sqrt{\frac{1}{n} \sum_{i=1}^n \left(Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right)^2} \\
& \times \sqrt{\frac{1}{n} \sum_{i=1}^n \left(\frac{1}{\prod_{\ell=s}^t \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^t e_\ell(H_{i,\ell}, a_\ell)} \right)^2} \\
& + \left(\frac{1}{\eta} \right) \sqrt{\frac{1}{n} \sum_{i=1}^n \left(Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right)^2} \\
& \times \sqrt{\frac{1}{n} \sum_{i=1}^n \left(\frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right)^2}
\end{aligned}$$

where the last inequality follows from Cauchy-Schwartz inequality and Assumption 2.3 (overlap condition). Maximizing over $\Pi_{(t+1):T}$ and taking the expectation of both sides yields:

$$\begin{aligned}
& E \left[\sup_{\pi^{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (C)}^{a,s,t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi^{(t+1):T}) \right| \right] \\
& \leq E \left[\sup_{\pi^{(t+1):T} \in \Pi_{(t+1):T}} \sqrt{\frac{1}{n} \sum_{i=1}^n \left(Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right)^2} \right. \\
& \times \left. \sqrt{\frac{1}{n} \sum_{i=1}^n \left(\frac{1}{\prod_{\ell=s}^t \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^t e_\ell(H_{i,\ell}, a_\ell)} \right)^2} \right] \\
& + \eta^{-1} E \left[\sup_{\pi^{(t+1):T} \in \Pi_{(t+1):T}} \sqrt{\frac{1}{n} \sum_{i=1}^n \left(Q_t^{\pi^{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi^{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right)^2} \right]
\end{aligned}$$

$$\begin{aligned}
& \times \sqrt{\frac{1}{n} \sum_{i=1}^n \left(\frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right)^2} \\
& \leq \sqrt{\frac{1}{n} \sum_{i=1}^n E \left[\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \left(Q_t^{\pi_{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right)^2 \right]} \\
& \times \sqrt{\frac{1}{n} \sum_{i=1}^n E \left[\left(\frac{1}{\prod_{\ell=s}^t \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^t e_\ell(H_{i,\ell}, a_\ell)} \right)^2 \right]} \\
& + \eta^{-1} \sqrt{\frac{1}{n} \sum_{i=1}^n E \left[\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \left(Q_t^{\pi_{(t+1):T}}(H_{i,t}, a_t) - \hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,t}, a_t) \right)^2 \right]} \\
& \times \sqrt{\frac{1}{n} \sum_{i=1}^n E \left[\left(\frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right)^2 \right]} \\
& = \sum_{k=1}^K \sqrt{E \left[\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \left(Q_t^{\pi_{(t+1):T}}(H_t, a_t) - \hat{Q}_t^{\pi_{(t+1):T}, -k}(H_{i,t}, a_t) \right)^2 \right]} \\
& \times \sqrt{E \left[\left(\frac{1}{\prod_{\ell=s}^t \hat{e}_\ell^{-k}(H_\ell, a_\ell)} - \frac{1}{\prod_{\ell=s}^t e_\ell(H_\ell, a_\ell)} \right)^2 \right]} \\
& + \eta^{-1} \sum_{k=1}^K \sqrt{E \left[\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \left(Q_t^{\pi_{(t+1):T}}(H_t, a_t) - \hat{Q}_t^{\pi_{(t+1):T}, -k}(H_t, a_t) \right)^2 \right]} \\
& \times \sqrt{E \left[\left(\frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_\ell, a_\ell)} - \frac{1}{\prod_{\ell=s}^{t-1} e_\ell(H_\ell, a_\ell)} \right)^2 \right]} \\
& = O(n^{-\tau/2}),
\end{aligned}$$

where the second inequality follows from Cauchy-Schwartz inequality and the last line follows from Assumption 4.1 (i). Then applying Markov's inequality leads to

$$\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (C)}^{a:s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| = O_P(n^{-\tau/2}), \quad (\text{B.17})$$

which proves equation (B.14).

Now let us consider $\tilde{S}_{s:t, (A2)}^{a:s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T})$. Note that

$$\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (A2)}^{a:s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right|$$

$$\begin{aligned}
&\leq \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \frac{1}{n} \sum_{i=1}^n \left| \frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right| \\
&\times \left| \hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,t}, a_t) - Q_t^{\pi_{(t+1):T}}(H_{i,t}, a_t) \right| \cdot \left| 1 - \frac{1\{A_{i,t} = a_t\}}{e_t(H_{i,t}, a_t)} \right| \\
&\leq \left(\frac{1}{\eta}\right) \sqrt{\frac{1}{n} \sum_{i=1}^n \left(\frac{1}{\prod_{\ell=s}^{t-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} - \frac{1}{\prod_{\ell=s}^{t-1} e_\ell(H_{i,\ell}, a_\ell)} \right)^2} \\
&\times \sqrt{\frac{1}{n} \sum_{i=1}^n \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \left(\hat{Q}_t^{\pi_{(t+1):T}, -k(i)}(H_{i,t}, a_t) - Q_t^{\pi_{(t+1):T}}(H_{i,t}, a_t) \right)^2},
\end{aligned}$$

where the last inequality follows from Assumption 2.3 (overlap condition) and Cauchy-Schwartz inequality. Then, by applying the same argument to derive (B.17), we obtain

$$\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (A2)}^{a,s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| = O_P(n^{-\tau/2}).$$

Combining this result with (B.15) leads to:

$$\begin{aligned}
&\sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (A)}^{a,s:t}(\pi_{s:t}^a, \pi_{s:t}^b, \pi_{(t+1):T}) \right| \\
&\leq \sum_{k=1}^K \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (A1)}^{a,s:t, k}(\pi_{s:t}) \right| + \sup_{\pi_{s:t}^a, \pi_{s:t}^b \in \Pi_{s:t}} \left| \tilde{S}_{s:t, (A2)}^{a,s:t}(\pi_{s:t}) \right| \\
&= O_P(n^{-\min\{1/2, \tau/2\}}).
\end{aligned}$$

This result proves equation (B.12).

Consequently, combining equations (B.11)–(B.14), we obtain the result (B.9).

We next consider the case that $t = T$. In this case, $\tilde{S}_{t:T}^{a,t:T}(\pi_{t:T}^a, \pi_{t:T}^b)$ is decomposed as

$$\tilde{S}_{t:T}^{a,t:T}(\pi_{t:T}^a, \pi_{t:T}^b) = \tilde{S}_{t:T, (A)}^{a,t:T}(\pi_{t:T}^a, \pi_{t:T}^b) + \tilde{S}_{t:T, (B)}^{a,t:T}(\pi_{t:T}^a, \pi_{t:T}^b) + \tilde{S}_{t:T, (C)}^{a,t:T}(\pi_{t:T}^a, \pi_{t:T}^b),$$

where

$$\begin{aligned}
\tilde{S}_{t:T, (A)}^{a,t:T}(\pi_{t:T}^a, \pi_{t:T}^b) &\equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{t:T}^a, \pi_{t:T}^b}^{a,t:T} \frac{\prod_{\ell=s}^{T-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{T-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \\
&\times \left(\hat{Q}_T^{-k(i)}(H_{i,T}, a_T) - Q_T(H_{i,T}, a_T) \right) \left(1 - \frac{1\{A_{i,T} = a_T\}}{e_T(H_{i,T}, a_T)} \right); \\
\tilde{S}_{t:T, (B)}^{a,t:T}(\pi_{t:T}^a, \pi_{t:T}^b) &\equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{t:T}^a, \pi_{t:T}^b}^{a,t:T} \frac{\prod_{\ell=s}^{T-1} 1\{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{T-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)}
\end{aligned}$$

$$\begin{aligned}
& \times (Y_{i,T} - Q_T(H_{i,T}, a_T)) \left(\frac{1 \{A_{i,T} = a_T\}}{\hat{e}_T^{-k(i)}(H_{i,T}, a_T)} - \frac{1 \{A_{i,T} = a_T\}}{e_T(H_{i,T}, a_T)} \right); \\
\tilde{S}_{t:T,(C)}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b) & \equiv \frac{1}{n} \sum_{i=1}^n G_{i, \pi_{t:T}^a, \pi_{t:T}^b}^{a_t:T} \frac{\prod_{\ell=s}^{T-1} 1 \{A_{i,\ell} = a_\ell\}}{\prod_{\ell=s}^{T-1} \hat{e}_\ell^{-k(i)}(H_{i,\ell}, a_\ell)} \\
& \times \left(Q_T(H_{i,T}, a_T) - \hat{Q}_T^{-k(i)}(H_{i,T}, a_T) \right) \left(\frac{1 \{A_{i,T} = a_T\}}{\hat{e}_T^{-k(i)}(H_{i,T}, a_T)} - \frac{1 \{A_{i,T} = a_T\}}{e_T(H_{i,T}, a_T)} \right).
\end{aligned}$$

The same arguments to derive the results (B.12)–(B.14) also show that

$$\begin{aligned}
\sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} \left| \tilde{S}_{t:T,(A)}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b) \right| &= O_P(n^{-1/2}); \\
\sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} \left| \tilde{S}_{t:T,(B)}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b) \right| &= O_P(n^{-1/2}); \\
\sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} \left| \tilde{S}_{t:T,(C)}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b) \right| &= O_P(n^{-\min\{1/2, \tau/2\}}).
\end{aligned}$$

Therefore,

$$\begin{aligned}
& \sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} \left| \tilde{S}_{t:T}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b) \right| \\
& \leq \sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} \left| \tilde{S}_{t:T,(A)}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b) \right| + \sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} \left| \tilde{S}_{t:T,(B)}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b) \right| \\
& + \sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} \left| \tilde{S}_{t:T,(C)}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b) \right| \\
& = O_P(n^{-\min\{1/2, \tau/2\}}),
\end{aligned}$$

which leads to the result (B.10). \square

We finally presents the proof of Lemma A.3.

Proof of Lemma A.3. From equation (B.8), it suffices to show that

$$\sum_{a_t \in \mathcal{A}_t} \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} |\tilde{S}_t^{a_t}(\pi_t^a, \pi_t^b, \pi_{(t+1):T})| = O_P(n^{-\min\{1/2, \tau/2\}}). \quad (\text{B.18})$$

Applying the result (B.9) in Lemma B.5 sequentially to the left-hand side of the above equation leads to

$$\sum_{a_t \in \mathcal{A}_t} \sup_{\pi_{(t+1):T} \in \Pi_{(t+1):T}} \sup_{\pi_t^a, \pi_t^b \in \Pi_t} |\tilde{S}_t^{a_t}(\pi_t^a, \pi_t^b, \pi_{(t+1):T})| \leq O_P(n^{-\min\{1/2, \tau/2\}})$$

$$+ \sum_{\underline{a}_{t:T} \in \underline{A}_{t:T}} \sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} |\tilde{S}_{t:T}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b)|.$$

The result (B.10) in Lemma B.5 shows that

$$\sup_{\pi_{t:T}^a, \pi_{t:T}^b \in \Pi_{t:T}} |\tilde{S}_{t:T}^{a_t:T}(\pi_{t:T}^a, \pi_{t:T}^b)| = O_P(n^{-\min\{1/2, \tau/2\}}).$$

Combining these results leads to the result (B.18), which ultimately proves the result in Lemma A.3. \square

C Proof of Theorem 5.1

This section presents the proof of Theorem 5.1.

Proof of Theorem 5.1. We begin by noting that the objective function $\widehat{W}^{AIPW}(\pi)$ can be expressed as $\widehat{W}^{AIPW}(\pi) = (1/n) \sum_{i=1}^n \widehat{V}_{i,1}(\pi)$, where $\widehat{V}_{i,1}(\cdot)$ is defined in Appendix A. A standard argument from statistical learning theory (e.g., Lugosi, 2002) leads to

$$\begin{aligned} R(\hat{\pi}^{AIPW}) &= \Delta_1(\pi^{*,opt}; \hat{\pi}^{AIPW}) \\ &\leq \Delta_1(\pi^{*,opt}; \hat{\pi}^{AIPW}) - \widehat{\Delta}_1(\pi^{*,opt}; \hat{\pi}^{AIPW}) \\ &\leq \sup_{\pi^a, \pi^b \in \Pi} |\Delta_1(\pi^a; \pi^b) - \widehat{\Delta}_1(\pi^a; \pi^b)| \\ &\leq \sup_{\pi^a, \pi^b \in \Pi} |\Delta_1(\pi^a; \pi^b) - \widetilde{\Delta}_1(\pi^a; \pi^b)| + \sup_{\pi^a, \pi^b \in \Pi} |\widehat{\Delta}_1(\pi^a; \pi^b) - \widetilde{\Delta}_1(\pi^a; \pi^b)|, \end{aligned} \quad (\text{C.1})$$

where the first inequality follows because $\hat{\pi}^{AIPW}$ maximizes $(1/n) \sum_{i=1}^n \widehat{V}_{i,1}(\pi)$ over Π ; hence, $\widehat{\Delta}_1(\pi^{*,opt}; \hat{\pi}^{AIPW}) \leq 0$.

Regarding $\sup_{\pi^a, \pi^b \in \Pi} |\Delta_1(\pi^a; \pi^b) - \widetilde{\Delta}_1(\pi^a; \pi^b)|$, under Assumptions 2.2, 2.3, and the result in Lemma 4.2, we can apply Lemma B.3 to obtain the following result: For any $\delta \in (0, 1)$, with probability at least $1 - 2\delta$,

$$\begin{aligned} \sup_{\pi^a, \pi^b \in \Pi} |\widetilde{\Delta}_1(\pi^a; \pi^b) - \Delta_1(\pi^a; \pi^b)| &\leq \left(54.4\sqrt{2}\kappa(\Pi) + 435.2 + \sqrt{2 \log \frac{1}{\delta}} \right) \sqrt{\frac{V^*}{n}} \\ &\quad + o\left(\frac{1}{\sqrt{n}}\right), \end{aligned} \quad (\text{C.2})$$

where $V^* \equiv \sup_{\pi^a, \pi^b \in \Pi} E \left[\left(\widetilde{\Gamma}_{i,1}^{\pi^a; T}(\pi^a) - \widetilde{\Gamma}_{i,1}^{\pi^b; T}(\pi^b) \right)^2 \right] < \infty$.

Regarding $\sup_{\pi^a, \pi^b \in \Pi} |\widehat{\Delta}_1(\pi^a; \pi^b) - \widetilde{\Delta}_1(\pi^a; \pi^b)|$, we have

$$\widehat{\Delta}_1(\pi^a; \pi^b) - \widetilde{\Delta}_1(\pi^a; \pi^b) = \sum_{\underline{a}_T \in \mathcal{A}_T} \widetilde{S}_{1:T}^{\underline{a}_T}(\pi^a, \pi^b).$$

Applying the result (B.10) in Lemma B.5 to $\widetilde{S}_{1:T}^{\underline{a}_T}(\pi^a, \pi^b)$ for each \underline{a}_T gives

$$\sup_{\pi^a, \pi^b \in \Pi} |\widehat{\Delta}_1(\pi^a; \pi^b) - \widetilde{\Delta}_1(\pi^a; \pi^b)| = O_p(n^{-\min\{1/2, \tau/2\}}). \quad (\text{C.3})$$

Combining the results (C.1)-(C.3) leads to the result (5). \square

Acknowledgements

I thank Toru Kitagawa, Aureo de Paula, Takahide Yanagi, and participants at various seminars and conferences for their comments and suggestions. I gratefully acknowledge financial support from JSPS KAKENHI Grant (number 22K20155) and ERC Grant (number 715940).

References

- ATHEY, S., J. TIBSHIRANI, AND S. WAGER (2019): “Generalized Random Forests,” *The Annals of Statistics*, 47, 1148–1178.
- ATHEY, S. AND S. WAGER (2021): “Policy Learning with Observational Data,” *Econometrica*, 89, 133–161.
- BHATTACHARYA, D. AND P. DUPAS (2012): “Inferring welfare maximizing treatment assignment under budget constraints,” *Journal of Econometrics*, 167, 168–196.
- BLUMLEIN, T., J. PERSSON, AND S. FEUERRIEGEL (2022): “Learning Optimal Dynamic Treatment Regimes Using Causal Tree Methods in Medicine,” in *Proceedings of the 7th Machine Learning for Healthcare Conference*, vol. 182, 146–171.
- BREIMAN, L., J. FRIEDMAN, C. STONE, AND R. OLSHEN (1984): *Classification and Regression Trees*, The Wadsworth and Brooks-Cole statistics-probability series, Taylor & Francis.

- CHAKRABORTY, B. AND S. A. MURPHY (2014): “Dynamic Treatment Regimes,” *Annual Review of Statistics and Its Application*, 1, 447–464.
- CHERNOZHUKOV, V., D. CHETVERIKOV, M. DEMIRER, E. DUFLO, C. HANSEN, W. NEWEY, AND J. ROBINS (2018): “Double/Debiased Machine Learning for Treatment and Structural Parameters,” *The Econometrics Journal*, 21, 1–68.
- DING, W. AND S. F. LEHRER (2010): “Estimating Treatment Effects from Contaminated Multiperiod Education Experiments: The Dynamic Impacts of Class Size Reductions,” *The Review of Economics and Statistics*, 92, 31–42.
- DOROUDI, S., P. S. THOMAS, AND E. BRUNSKILL (2018): “Importance Sampling for Fair Policy Selection,” in *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 5239–5243.
- ERTEFAIE, A., J. R. MCKAY, D. OSLIN, AND R. L. STRAWDERMAN (2021): “Robust Q-learning,” *Journal of the American Statistical Association*, 116, 368–381.
- FONTENEAU, R., S. A. MURPHY, L. WEHENKEL, AND D. ERNST (2013): “Batch Mode Reinforcement Learning based on the Synthesis of Artificial Trajectories,” *Annals of Operations Research*, 208, 383–416.
- HAN, S. (2023): “Optimal dynamic treatment regimes and partial welfare ordering,” *Journal of the American Statistical Association*, 1–11.
- HECKMAN, J. J., J. E. HUMPHRIES, AND G. VERAMENDI (2016): “Dynamic treatment effects,” *Journal of Econometrics*, 191, 276–292.
- HECKMAN, J. J. AND S. NAVARRO (2007): “Dynamic discrete choice and dynamic treatment effects,” *Journal of Econometrics*, 136, 341–396.
- HERNÁN, M. A., B. BRUMBACK, AND J. M. ROBINS (2001): “Marginal Structural Models to Estimate the Joint Causal Effect of Nonrandomized Treatments,” *Journal of the American Statistical Association*, 96, 440–448.
- HIRANO, K. AND J. PORTER (2009): “Asymptotics for statistical treatment rules,” *Econometrica*, 77, 1683–1701.

- IDA, T., T. ISHIHARA, K. ITO, D. KIDO, T. KITAGAWA, S. SAKAGUCHI, AND S. SASAKI (2024): “Dynamic Targeting: Experimental Evidence from Energy Rebate Programs,” *NBER Working Paper: 32561*.
- JIANG, N. AND L. LI (2016): “Doubly Robust Off-policy Value Evaluation for Reinforcement Learning,” in *Proceedings of the 33rd International Conference on Machine Learning*, PMLR, 652–661.
- KALLUS, N. AND M. UEHARA (2020): “Double Reinforcement Learning for Efficient Off-Policy Evaluation in Markov Decision Processes,” *Journal of Machine Learning Research*, 21, 6742–6804.
- KITAGAWA, T., S. SAKAGUCHI, AND A. TETENOV (2023): “Constrained Classification and Policy Learning,” *arXiv preprint arXiv:2106.12886*.
- KITAGAWA, T. AND A. TETENOV (2018): “Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice,” *Econometrica*, 86, 591–616.
- KOLSRUD, J., C. LANDAIS, P. NILSSON, AND J. SPINNEWIJN (2018): “The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden,” *American Economic Review*, 108, 985–1033.
- KOSOROK, M. R. AND E. B. LABER (2019): “Precision Medicine,” *Annual Review of Statistics and Its Application*, 6, 263–286.
- KRUEGER, A. B. (1999): “Experimental Estimates of Education Production Functions,” *The Quarterly Journal of Economics*, 114, 497–532.
- LABER, E. B., D. J. LIZOTTE, M. QIAN, W. E. PELHAM, AND S. A. MURPHY (2014): “Dynamic Treatment Regimes: Technical Challenges and Applications,” *Electronic Journal of Statistics*, 8, 1225–1272.
- LABER, E. B. AND Y.-Q. ZHAO (2015): “Tree-Based Methods for Individualized Treatment Regimes,” *Biometrika*, 102, 501–514.
- LE, H., C. VOLOSHIN, AND Y. YUE (2019): “Batch Policy Learning under Constraints,” in *International Conference on Machine Learning*, PMLR, 3703–3712.
- LECHNER, M. (2009): “Sequential Causal Models for the Evaluation of Labor Market Programs,” *Journal of Business & Economic Statistics*, 27, 71–83.

- LI, Z., J. CHEN, E. LABER, F. LIU, AND R. BAUMGARTNER (2023): “Optimal Treatment Regimes: A Review and Empirical Comparison,” *International Statistical Review*, 91, 427–463.
- LIU, X. (2023): “Dynamic Coupon Targeting Using Batch Deep Reinforcement Learning: An Application to Livestream Shopping,” *Marketing Science*, 42, 637–658.
- LUGOSI, G. (2002): “Pattern Classification and Learning Theory,” in *Principles of Non-parametric Learning*, ed. by L. Györfi, Vienna: Springer, 1–56.
- MANSKI, C. F. (2004): “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, 72, 1221–1246.
- MBAKOP, E. AND M. TABORD-MEEHAN (2021): “Model selection for treatment choice: Penalized welfare maximization,” *Econometrica*, 89, 825–848.
- MEYER, B. D. (1995): “Lessons from the U.S. Unemployment Insurance Experiments,” *Journal of Economic Literature*, 33, 91–131.
- MOODIE, E., B. CHAKRABORTY, AND M. S. KRAMER (2012): “Q-learning for Estimating Optimal Dynamic Treatment Rules from Observational Data,” *Canadian Journal of Statistics*, 40, 629–645.
- MUNOS, R. AND C. SZEPESVÁRI (2008): “Finite-Time Bounds for Fitted Value Iteration,” *Journal of Machine Learning Research*, 9, 815–857.
- MURPHY, S. A. (2003): “Optimal Dynamic Treatment Regimes,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65, 331–355.
- (2005): “A Generalization Error for Q-learning,” *Journal of Machine Learning Research*, 6, 1073–1097.
- NIE, X., E. BRUNSKILL, AND S. WAGER (2021): “Learning When-to-Treat Policies,” *Journal of the American Statistical Association*, 116, 392–409.
- PELHAM JR, W. E., G. A. FABIANO, J. G. WAXMONSKY, A. R. GREINER, E. M. GNAGY, W. E. PELHAM III, S. COXE, J. VERLEY, I. BHATIA, K. HART, ET AL. (2016): “Treatment Sequencing for Childhood ADHD: A Multiple-Randomization Study of Adaptive Medication and Behavioral Interventions,” *Journal of Clinical Child & Adolescent Psychology*, 45, 396–415.

- QIAN, M. AND S. A. MURPHY (2011): “Performance Guarantees for Individualized Treatment Rules,” *The Annals of Statistics*, 39, 1180.
- ROBINS, J. M. (1986): “A New Approach to Causal Inference in Mortality Studies With a Sustained Exposure Period—Application to Control of the Healthy Worker Survivor Effect,” *Mathematical Modelling*, 7, 1393–1512.
- (1997): “Causal Inference From Complex Longitudinal Data in Latent Variable Modeling and Applications to Causality,” in *Lecture Notes in Statistics*, ed. by M. Berkane, New York: Springer, 69–117.
- (2004): “Optimal Structural Nested Models for Optimal Sequential Decisions,” in *Proceedings of the Second Seattle Symposium in Biostatistics. Lecture Notes in Statistics*, ed. by D. Y. Lin and P. J. Heagerty, New York: Springer, 189–326.
- RODRÍGUEZ, J., F. SALTIEL, AND S. URZÚA (2022): “Dynamic Treatment Effects of Job Training,” *Journal of Applied Econometrics*, 37, 242–269.
- SAKAGUCHI, S. (2024): “Estimation of Optimal Dynamic Treatment Assignment Rules Under Policy Constraints,” *arXiv preprint arXiv:2106.05031*.
- SCHICK, A. (1986): “On Asymptotically Efficient Estimation in Semiparametric Models,” *The Annals of Statistics*, 1139–1151.
- STOYE, J. (2009): “Minimax regret treatment choice with finite samples,” *Journal of Econometrics*, 151, 70–81.
- (2012): “Minimax regret treatment choice with covariates or with limited validity of experiments,” *Journal of Econometrics*, 166, 138–156.
- SUN, Y. AND L. WANG (2021): “Stochastic Tree Search for Estimating Optimal Dynamic Treatment Regimes,” *Journal of the American Statistical Association*, 116, 421–432.
- TAO, Y., L. WANG, AND D. ALMIRALL (2018): “Tree-Based Reinforcement Learning for Estimating Optimal Dynamic Treatment Regimes,” *The Annals of Applied Statistics*, 12, 1914–1938.
- TETENOV, A. (2012): “Statistical treatment choice based on asymmetric minimax regret criteria,” *Journal of Econometrics*, 166, 157–165.

- THOMAS, P. AND E. BRUNSKILL (2016): “Data-Efficient Off-Policy Policy Evaluation for Reinforcement Learning,” in *Proceedings of the 33rd International Conference on Machine Learning*, PMLR, 2139–2148.
- TSIATIS, A. A., M. DAVIDIAN, S. T. HOLLOWAY, AND E. B. LABER (2019): *Dynamic Treatment Regimes: Statistical Methods for Precision Medicine*, Chapman and Hall/CRC.
- VIVIANO, D. (2024): “Policy targeting under network interference,” *Review of Economic Studies*, rdae041.
- WALLACE, M. P. AND E. E. MOODIE (2015): “Doubly-Robust Dynamic Treatment Regimen Estimation via Weighted Least Squares,” *Biometrics*, 71, 636–644.
- WANG, L., A. ROTNITZKY, X. LIN, R. E. MILLIKAN, AND P. F. THALL (2012): “Evaluation of Viable Dynamic Treatment Regimes in a Sequentially Randomized Trial of Advanced Prostate Cancer,” *Journal of the American Statistical Association*, 107, 493–508.
- WATKINS, C. J. AND P. DAYAN (1992): “Q-learning,” *Machine learning*, 8, 279–292.
- WORD, E., J. JOHNSTON, H. P. BAIN, B. D. FULTON, J. B. ZAHARIES, M. N. LINTZ, C. M. ACHILLES, J. FOLGER, AND C. BREDI (1990): “The State of Tennessee’s Student/Teacher Achievement Ratio (STAR) Project: Technical Report (1985-1990).” Tennessee State Department of Education.
- ZHANG, B., A. A. TSIATIS, E. B. LABER, AND M. DAVIDIAN (2013): “Robust Estimation of Optimal Dynamic Treatment Regimes for Sequential Treatment Decisions,” *Biometrika*, 100, 681–694.
- ZHANG, Y., E. B. LABER, M. DAVIDIAN, AND A. A. TSIATIS (2018): “Interpretable Dynamic Treatment Regimes,” *Journal of the American Statistical Association*, 113, 1541–1549.
- ZHAO, Y., D. ZENG, A. J. RUSH, AND M. R. KOSOROK (2012): “Estimating Individualized Treatment Rules Using Outcome Weighted Learning,” *Journal of the American Statistical Association*, 107, 1106–1118.

ZHAO, Y. Q., D. ZENG, E. B. LABER, AND M. R. KOSOROK (2015): “New Statistical Learning Methods for Estimating Optimal Dynamic Treatment Regimes,” *Journal of the American Statistical Association*, 110, 583–598.

ZHOU, N., L. WANG, AND D. ALMIRALL (2023a): “Estimating tree-based dynamic treatment regimes using observational data with restricted treatment sequences,” *Biometrics*, 79, 2260–2271.

ZHOU, Z., S. ATHEY, AND S. WAGER (2023b): “Offline Multi-Action Policy Learning: Generalization and Optimization,” *Operations Research*, 71, 148–183.