

Demirer, Mert; Syrgkanis, Vasilis; Lewis, Gregory; Chernozhukov, Victor

Working Paper

Semi-parametric efficient policy learning with continuous actions

cemmap working paper, No. CWP34/19

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Demirer, Mert; Syrgkanis, Vasilis; Lewis, Gregory; Chernozhukov, Victor (2019) : Semi-parametric efficient policy learning with continuous actions, cemmap working paper, No. CWP34/19, Centre for Microdata Methods and Practice (cemmap), London, <https://doi.org/10.1920/wp.cem.2019.3419>

This Version is available at:

<https://hdl.handle.net/10419/211127>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Semi-Parametric Efficient Policy Learning with Continuous Actions

Mert Demirer
Vasilis Syrgkanis
Greg Lewis
Victor Chernozhukov

The Institute for Fiscal Studies
Department of Economics,
UCL

cemmap working paper CWP34/19

Semi-Parametric Efficient Policy Learning with Continuous Actions

Mert Demirer

MIT

mdemirer@mit.edu

Vasilis Syrgkanis

Microsoft Research

vasy@microsoft.com

Greg Lewis

Microsoft Research

glewis@microsoft.com

Victor Chernozhukov

MIT

vchern@mit.edu

Abstract

We consider off-policy evaluation and optimization with continuous action spaces. We focus on observational data where the data collection policy is unknown and needs to be estimated. We take a semi-parametric approach where the value function takes a known parametric form in the treatment, but we are agnostic on how it depends on the observed contexts. We propose a doubly robust off-policy estimate for this setting and show that off-policy optimization based on this estimate is robust to estimation errors of the policy function or the regression model. Our results also apply if the model does not satisfy our semi-parametric form, but rather we measure regret in terms of the best projection of the true value function to this functional space. Our work extends prior approaches of policy optimization from observational data that only considered discrete actions. We provide an experimental evaluation of our method in a synthetic data example motivated by optimal personalized pricing and costly resource allocation.

1 Introduction

We consider off-policy evaluation and optimization with continuous action spaces from observational data, where the data collection (logging) policy is unknown. We take a semi-parametric approach where we assume that the value function takes a known parametric form in the treatment, but we are agnostic on how it depends on the observed contexts/features. In particular, we assume that:

$$V(a, z) = \langle \theta_0(z), \phi(a, z) \rangle \quad (1)$$

for some known feature functions ϕ but unknown functions θ_0 . We assume that we are given a set of n observational data points (x_1, \dots, x_n) that consist of i.i.d copies of the random vector $x = (y, a, z) \in \mathcal{Y} \times \mathcal{A} \times \mathcal{Z}$, such that $\mathbb{E}[y \mid a, z] = V(a, z)$.¹

Our goal is to estimate a policy $\hat{\pi} : \mathcal{Z} \rightarrow \mathcal{A}$ from a space of policies Π that achieves good regret:

$$\sup_{\pi \in \Pi} \mathbb{E}[V(\pi(z), z)] - \mathbb{E}[V(\hat{\pi}(z), z)] \leq R(\Pi, n) \quad (2)$$

for some regret rate that depends on the policy space Π and the sample size n .

The semi-parametric value assumption allows us to formulate a doubly robust estimate V_{DR} of the value function, from the observational data, which depends on first stage regression estimates of the

¹In most of the paper, we can allow for the case where z is endogenous, in the sense that $\mathbb{E}[y \mid a, z] = V(a, z) + f_0(z)$. In other words, the noise in the random variable y can be potentially correlated with z . However, we assume that conditional on z , there is no remaining endogeneity in the choice of the action in our data. The latter is typically referred to as conditional ignorability/exogeneity [11].

coefficients $\theta_0(z)$ and the conditional covariance of the features $\Sigma_0(z) = \mathbb{E}[\phi(a, z)\phi(a, z)^T \mid z]$. The latter is the analogue of the propensity function when actions are discrete. Our estimate is doubly robust in that it is unbiased if either θ_0 or Σ_0 is correct. Then we optimize this estimate:

$$\hat{\pi} = \sup_{\pi \in \Pi} V_{DR}(\pi) \quad (3)$$

Main contributions. We show that the double robustness property implies that our objective function satisfies a Neyman orthogonality criterion, which in turn implies that our regret rates depend only in a second order manner on the estimation errors on the first stage regression estimates of the functions θ_0, Σ_0 . Moreover, we prove a regret rate whose leading term depends on the variance of the difference of our estimated value between any two policy values within a “small regret-slice” and on the entropy integral of the policy space. We achieve this with a computationally efficient variant of the empirical risk minimization (ERM) algorithm (of independent interest) that uses a validation set to construct a preliminary policy and use it to regularize the policy computed on the training set. Hence, we manage to achieve variance-based regret bounds without the need for variance or moment penalization [15, 20, 9] used in prior work and which can render a computationally tractable policy learning problem, non-convex. We also show that the asymptotic variance of our off-policy estimate (which governs the scale of the leading regret term) is asymptotic minimax optimal, in the sense that it achieves the semi-parametric efficiency lower bound.

Robustness to mis-specification. Notably, our approach provides meaningful guarantees even when our semi-parametric value function assumption is violated. Suppose that the true value function does not take the form of Equation (1), but rather takes some other form $V_0(a, z)$. Then one can consider the projection of this value function onto the forms of Equation (1), as:

$$\theta_p(z) = \arg \inf_{\theta} \mathbb{E} [(V_0(a, z) - \langle \theta(z), \phi(a, z) \rangle)^2 \mid z] \quad (4)$$

where the expectation is taken over the distribution of observed data. Then our approach takes the interpretation of achieving good regret bounds with respect to this best linear semi-parametric approximation. This is an alternative to the kernel smoothing approximation proposed by [20] in contextual bandit setting, as a regret target, and related to [12]. If there is some rough domain knowledge on the form of how the action affects the reward, then our semi-parametric approximate target should achieve better performance when the dimension of the action space is large, as the bias of kernel methods will typically incur an exponential in the dimension bias.

Double robustness. In cases where the collection policy is known, our doubly robust approach can be used for variance reduction via fitting first stage regression estimates to the policy value, whilst maintaining unbiasedness. Thus we can apply our approach to improve regret in the counterfactual risk minimization framework [20], [12] and as a variance reduction method in contextual bandit algorithms with continuous actions [20].

Related Literature. Our work builds on the recent work at the intersection of semi-parametric inference and policy learning from observational data. The important work of [1] analyzes the binary treatment setting and also takes a doubly robust approach so as to obtain regret bounds whose leading term depends on the semi-parametric efficient variance and the entropy integral and which is robust to first stage estimation errors. Hence, our results are a direct generalization of their work to arbitrary continuous actions spaces, subject to our semi-parametric value assumption. In fact we show formally in the Appendix how one can recover the main result of [1], from our main regret bound. In turn our work builds on a long line of work on policy learning and counterfactual risk minimization [17, 26, 27, 1, 13, 28, 2, 7, 20, 12, 14]. Notably, the work of [28] extends the work of [1] to many discrete actions, but only proves a second moment based regret bound, which can be much larger than the variance. Our setting also subsumes the setting of many discrete actions and hence our regularized ERM offers an improvement over the rates in [28]. [9] formulates a general framework of statistical learning with a nuisance component. Our method falls into this framework and we build upon some of the results in [9]. However, for the case of policy learning the implications of [9] provide a variance based regret only when invoking second moment penalization, which can be intractable. We side-step this need and provide a computationally efficient alternative. Finally, most of the work on policy learning in machine learning assumes that the current policy (equiv. $\Sigma_0(z)$) is known. Hence, double robustness is used mostly as a variance reduction technique. Even for

this literature, as we discuss above, our method can be seen an alternative of recent work on policy learning with continuous actions [12, 14] that makes use of non-parametric kernel methods.

Our work also connects to the semi-parametric estimation literature in econometrics and statistics. Our model is an extension of the partially linear model which has been extensively studied in the econometrics [8, 19]. By considering context-specific coefficients (random coefficients) and modeling a value function that is non-linear in treatment, we substantially extend the partially linear model. [24, 10] studied a special case of our model where output is linearly dependent on treatment given context, with the aim of estimating the average treatment effect. [10] constructed the doubly robust estimator and showed its semi-parametric efficiency under the linear-in-treatment assumption. We extend their results to a more general functional form and use the double-robustness property and semi-parametric efficiency for policy evaluation and optimization rather than treatment effect estimation. Our work is also connected to the recent and rapidly growing literature on the orthogonal/locally robust/debiased estimation literature [5, 6, 21].

2 Orthogonal Off-Policy Evaluation and Optimization

Let $\hat{\theta}$ be a first stage estimate of $\theta_0(z)$, which can be obtained by minimizing the square loss:

$$\hat{\theta} = \arg \inf_{\theta \in \Theta} \mathbb{E}_n \left[(y - \langle \theta(z), \phi(a, z) \rangle)^2 \right] \quad (5)$$

where Θ is an appropriate parameter space for the parameters $\theta(z)$. Let $\Sigma_0(z)$ denote the conditional covariance matrix:

$$\Sigma_0(z) = \mathbb{E}[\phi(a, z) \phi(a, z)^T \mid z]$$

This is the analogue of the propensity model in discrete treatment settings. An estimate $\hat{\Sigma}(z)$ can be obtained by running a multi-task regression problem for each entry to the matrix, i.e.:

$$\hat{\Sigma}_{ij} = \arg \inf_{\Sigma_{ij} \in \mathcal{S}_{ij}} \mathbb{E} \left[(\phi_i(a, z) \phi_j(a, z) - \Sigma_{ij}(z))^2 \right] \quad (6)$$

where \mathcal{S}_{ij} is some appropriate hypothesis space for these regressions. Finally, the doubly robust estimate of the off-policy value takes the form:

$$V_{DR}(\pi) = \mathbb{E}_n [v_{DR}(y, a, z; \pi)] \quad (7)$$

where:

$$v_{DR}(y, a, z; \pi) = \langle \theta_{DR}(y, a, z), \phi(\pi(z), z) \rangle \quad (8)$$

$$\theta_{DR}(y, a, z) = \hat{\theta}(z) + \hat{\Sigma}(z)^{-1} \phi(a, z) (y - \langle \hat{\theta}(z), \phi(a, z) \rangle) \quad (9)$$

The quantity $\theta_{DR}(y, a, z)$ can be viewed as an estimate of $\theta_0(z)$, based on a single observation. In fact, if the matrix $\hat{\Sigma}$ was equal to Σ_0 , then one can see that $\theta_{DR}(y, a, z)$ is an unbiased estimate of $\theta_0(z)$. Our estimate v_{DR} also satisfies a doubly robust property, i.e. it is correct if either $\hat{\theta}$ is unbiased or $\hat{\Sigma}^{-1}$ is unbiased (see Appendix E for a formal statement). Finally, we will denote with $\theta_{DR}^0(y, a, z)$ the version of θ_{DR} , where the nuisance quantities θ and Σ are replaced by their true values, and correspondingly define $v_{DR}^0(y, a, z; \pi)$. We perform policy optimization based on this doubly robust estimate:

$$\hat{\pi} = \arg \max_{\pi \in \Pi} V_{DR}(\pi) \quad (10)$$

Moreover, we let π_*^0 be the optimal policy:

$$\pi_*^0 = \arg \max_{\pi \in \Pi} V(\pi) \quad (11)$$

Remark 1 (Multi-Action Policy Learning). *A special case of our setup is the setting where the number of actions is finitely many. This can be encoded as $a \in \{e_1, \dots, e_n\}$ and $\phi(a, z) = a$. In that case, observe that the covariance matrix becomes a diagonal matrix: $\Sigma_0(z) = \text{diag}(p_1(z), \dots, p_n(z))$, with $p_i(z) = \Pr[a = e_i \mid z]$. In this case, we simply recover the standard doubly robust estimate that combines the direct regression part with the inverse propensity weights part, i.e.:*

$$\theta_{DR,i}(y, a, z) = \hat{\theta}_i(z) + \frac{1}{\hat{p}_i(z)} 1\{a = e_i\} (y - \hat{\theta}_i(z))$$

Thus our estimator is an extension of the doubly robust estimate from discrete to continuous actions.

Remark 2 (Finitely Many Possible Actions: Linear Contextual Bandits). *Another interesting special case of our approach is a generalization of the linear contextual bandit setting. In particular, suppose that there is only a finite (albeit potentially large) set of $N > p$ possible actions $A = \{a_1, \dots, a_N\}$ and $a_i \in \mathbb{R}^p$. However, unlike the multi-action setting, where these actions are the orthonormal basis vectors, in this setting, each action $a_i \in A$, maps to a feature vector $\phi_i(z) := \phi(a_i, z)$. Then the reward y that we observe satisfies $\mathbb{E}[y | z, a] = \langle \theta(z), \phi(a, z) \rangle$. This is a generalization of the linear contextual bandit setting, in which the coefficient vector $\theta(z)$ is a constant parameter θ as opposed to varying with z . In this case observe that: $\Sigma_0(z) = \sum_{i=1}^N p_i(z) \phi_i(z) \phi_i(z)^T = U D U^T$, i.e. it is the sum of N rank one matrices where $D = \text{diag}(p_1(z), \dots, p_N(z))$, $p_i(z) = \Pr[a = a_i | z]$ and $U = [\phi_1(z), \dots, \phi_N(z)]$. The doubly robust estimate of the parameter takes the form:*

$$\theta_{DR}(y, a, z) = \hat{\theta}(z) + (U D U^T)^{-1} \phi(a, z) (y - \langle \hat{\theta}(z), \phi(a, z) \rangle)$$

This approach leverages the functional form assumption to get an estimate that avoids a large variance that depends on the number of actions N but rather mostly depends on the number of parameters p . This is achieved by sharing reward information across actions.

Remark 3 (Linear-in-Treatment Value). *Consider the case where the value is linear in the action $\phi(a, z) = a \in \mathbb{R}^p$. In this case observe that: $\Sigma_0(z) = \text{Var}(a | z)$. For instance, suppose that we assume that experimentation is independent across actions in the observed data. Then $\Sigma_0(z) = \text{diag}(\sigma_1^2(z), \dots, \sigma_p^2(z))$, where $\sigma_i^2 = \text{Var}(a_i | z)$. Then the doubly robust estimate of the parameter takes the form:*

$$\theta_{DR,i}(y, a, z) = \hat{\theta}_i(z) + \frac{a_i}{\hat{\sigma}_i^2(z)} (y - \langle \hat{\theta}(z), a \rangle) \quad (12)$$

3 Theoretical Analysis

Our main regret bounds are derived for a slight variation of the ERM algorithm that we presented in the preliminary section. In particular, we crucially need to augment the ERM algorithm with a “validation” step, where we split our data into a training and validation step, and we restrict attention to policies that achieve small regret on the training data, while still maintaining small regret on the validation set. This extra modification enabled us to prove variance based regret bounds and is reminiscent of standard approaches in machine learning, like k -fold cross-validation and early stopping, hence could be of independent interest.

Algorithm 1: Out-of-Sample Regularized ERM with Nuisance Estimates

- 1 The inputs are given by the sample of data $S = (x_1, \dots, x_n)$, which we randomly split in two parts S_1, S_2 . Moreover, we randomly split S_2 into validation and training samples S_2^v and S_2^t .
- 2 Estimate the nuisance functions $\hat{\theta}(z)$ and $\hat{\Sigma}(z)$ using Equations (5) and (6) on S_1 .
- 3 Use the output of Step 2 to construct the doubly robust moment in Equation (7) on S_2^v . Run ERM given in Equation (10) over policy space Π_1 on S_2^v to learn a policy function π_1 .
- 4 Use the output of Step 3 to construct a function class Π_2 defined as

$$\Pi_2 = \{\pi \in \Pi : \mathbb{E}_{S^v}[v_{DR}(y, a, z; \pi_1) - v_{DR}(y, a, z; \pi)] \leq \mu_n\}$$

for some μ_n and \mathbb{E}_{S^v} denotes the empirical expectation over the validation sample.

- 5 Use the output of Step 1 to construct the doubly robust moment in Equation (7) on S_2^t . Run a constrained ERM on S_2^t over Π_2 .
-

We note that we present our theoretical results for the simpler case where the nuisance estimates are trained on a separate split of the data. However, our results qualitatively extend to the case where we use the cross-fitting idea of [5] (i.e. train a model on one half and predict on the other and vice versa).

Regret bound. To show the properties of this algorithm, we first show that the regret of the doubly robust algorithm is impacted in a second order manner by the errors in the first stage estimates. We will also make the following preliminary definitions. For any function f we denote with $\|f\|_2 = \sqrt{\mathbb{E}[f(x)^2]}$, the standard L^2 norm and with $\|f\|_{2,n} = \sqrt{\mathbb{E}_n[f(x)^2]}$ its empirical analogue. Furthermore, we define the empirical entropy of a function class $H_2(\epsilon, \mathcal{F}, n)$ as the largest value, over the choice of n samples, of the logarithm of the size of the smallest empirical ϵ -cover of \mathcal{F} on

the samples with respect to the $\|\cdot\|_{2,n}$ norm. Finally, we consider the empirical entropy integral:

$$\kappa(r, \mathcal{F}) = \inf_{\alpha \geq 0} \left\{ 4\alpha + 10 \int_{\alpha}^r \sqrt{\frac{\mathcal{H}_2(\epsilon, \mathcal{F}, n)}{n}} d\epsilon \right\}, \quad (13)$$

Our statistical learning problem corresponds to learning over the function space:

$$\mathcal{F}_{\Pi} = \{v_{DR}(\cdot; \pi) : \pi \in \Pi\} \quad (14)$$

where the data is $x = (y, a, z)$. We will also make a very benign assumption on the entropy integral:

ASSUMPTION 1. *The function class \mathcal{F}_{Π} satisfies that for any constant r , $\kappa(r, \mathcal{F}) \rightarrow 0$ as $n \rightarrow \infty$.*

Theorem 1 (Variance-Based Oracle Policy Regret). *Suppose that the nuisance estimates satisfy that their mean squared error is upper bounded w.p. $1 - \delta/2$ by $h_{n,\delta}^2$, i.e. w.p. $1 - \delta/2$ over the randomness of the nuisance sample:*

$$\max \left\{ \mathbb{E}[(\hat{\theta}(z) - \theta_0(z))^2], \mathbb{E}[\|\hat{\Sigma}(z) - \Sigma_0(z)\|_{Fro}^2] \right\} \leq h_{n,\delta}^2 \quad (15)$$

Let $r = \sup_{\pi \in \Pi} \sqrt{\mathbb{E}[v_{DR}(z; \pi)^2]}$ and $\mu_n = \Theta \left(\kappa(r, \mathcal{F}_{\Pi}) + r \sqrt{\frac{\log(1/\delta)}{n}} \right)$. Moreover, let

$$\Pi_*(\epsilon) = \{\pi \in \Pi : V(\pi_*^0) - V(\pi) \leq \epsilon\}, \quad (16)$$

denote an ϵ -regret slice of the policy space. Let $\epsilon_n = O(\mu_n + h_{n,\delta}^2)$ and

$$V_2^0 = \sup_{\pi, \pi' \in \Pi_*(\epsilon_n)} \text{Var}(v_{DR}^0(x; \pi) - v_{DR}^0(x; \pi')) \quad (17)$$

denote the variance of the difference between any two policies in an ϵ_n -regret slice, evaluated at the true nuisance quantities. Then the policy π_2 returned by the out-of-sample regularized ERM, satisfies w.p. $1 - \delta$ over the randomness of S :

$$V(\pi_*^0) - V(\pi_2) = O \left(\kappa(\sqrt{V_2^0}, \mathcal{F}_{\Pi}) + \sqrt{\frac{V_2^0 \log(1/\delta)}{n}} + h_{n,\delta}^2 \right) \quad (18)$$

Expected regret is $O \left(\kappa(\sqrt{V_2^0}, \mathcal{F}_{\Pi}) + \sqrt{\frac{V_2^0}{n}} + h_n^2 \right)$, with h_n^2 is expected MSE of nuisance functions.

We provide a proof of this Theorem in Appendix B. The regret result contains two main contributions: 1) first the impact of the nuisance estimation error is of second order (i.e. $h_{n,\delta}^2$ instead of $h_{n,\delta}$), 2) the leading regret term depends on the variance of small-regret policy differences and the entropy integral of the policy space. The first property stems from the Neyman orthogonality property of the doubly robust estimate of the policy. The second property stems from the out-of-sample regularization step that we added to the ERM algorithm. Typically, we will have $h_{n,\delta}^2 = o(1/\sqrt{n})$ and thereby this term is of lower order than the leading term. Moreover, for many policy spaces $\kappa(0, \mathcal{F}_{\Pi}) = 0$, in which case we see that if the setting satisfies a “margin” condition (i.e. the best policy is better by a constant Δ margin), then eventually the variance of small regret policies is 0, since it only contains the best policy. In that case, our bound leads to fast rates of $\log(n)/n$ as opposed to $1/\sqrt{n}$, since the leading term vanishes (similar to the $\log(n)/n$ achieved in bandit problems with such a margin condition).

Dependence on the quantity V_2^0 is quite intuitive: if two policies have almost equivalent regret up to a μ_n rate, then it will be very easy to be misled among them if one has much higher variance than the other. For some classes of problems, the above also implies a regret rate that only depends on the variance of the optimal policy (e.g. when all policies with low regret have a variance that is not much larger than the variance of the optimal policy). In Appendix F we show that the latter is always the case for the setting of binary treatment studied in [1] and therefore applying our main result, we recover exactly their bound.

Semi-parametric efficient variance. Our regret bound depends on the variance of our doubly robust estimate of the value function. One then wonders if there are other estimates of the value function that could achieve better variance than $V_{DR}(\pi)$. However, we show that at least asymptotically and without further assumptions on the functions $\theta_0(z)$ and $\Sigma_0(z)$, this cannot be the case. In particular,

we show that our estimator achieves what is known as the semi-parametric efficient variance limit for our setting. More importantly for our regret result, this is also true for the semiparametric efficient variance of the policy differences. This is the case in our main setup; where the model is mis-specified and only a projection of the true value; and even if we assume that our model is correct, but make the extra assumption of homoskedasticity, i.e., the conditional variance of residuals of outcomes y do not depend on (a, z) .

Theorem 2 (Semi-parametric Efficiency). *If the model is mis-specified, i.e., $V_0(a, z) \neq V(a, z)$ the asymptotic variance of $V_{DR}(\pi)$ is equal to the semi-parametric efficiency bound for the policy value $\langle \theta_p(z), \pi(z) \rangle$ defined in Equation (4). If the model is correctly specified, $V_{DR}(\pi)$ is semi-parametrically efficient under homoskedasticity, i.e. $\mathbb{E}[(y - V(a, z))^2 | a, z] = \mathbb{E}[(y - V(a, z))^2]$.*

We provide a proof for the value function, but this result also extends to the difference of values. We conclude the section by providing concrete examples of rates for policy classes of interest.

Example 1 (VC Policies). *As a concrete example, consider the case when the class \mathcal{F}_Π is a VC-subgraph class of VC dimension d (e.g. the policy space has small VC-dimension or pseudo-dimension), and let $S = \mathbb{E}[\sup_\pi v_{DR}(x; \pi)^2]$. Then Theorem 2.6.7 of [22] shows that: $\mathcal{H}_2(\epsilon, \mathcal{F}_\Pi, n) = O(d(1 + \log(S/\epsilon)))$ (see also discussion in Appendix F). This implies that*

$$\kappa(r, \mathcal{F}_\Pi) = O\left(\int_0^r \sqrt{d(1 + \log(S/\epsilon))} d\epsilon\right) = O\left(r\sqrt{d}\sqrt{1 + \log(S/r)}\right).$$

Hence, we can conclude that regret is $O\left(\sqrt{V_2^0(1 + \log(S/V_2^0))}\sqrt{\frac{d}{n}} + \sqrt{\frac{V_2^0 \log(1/\delta)}{n}} + h_{n,\delta}^2\right)$. For the case of binary action policies (as we discuss in Appendix F) this result recovers the result of [1] up to constants and extends it to arbitrary action spaces and VC-subgraph policies.

Example 2 (High-Dimensional Index Policies). *As an example, we consider the class of policies, characterized by a constant number of ℓ_1 or ℓ_0 -bounded linear indices:*

$$\Pi_1 = \{z \rightarrow \Gamma(\langle \beta_1, z \rangle, \dots, \langle \beta_d, z \rangle) : \beta_i \in \mathbb{R}^p, \|\beta_i\|_1 \leq s\} \quad (19)$$

where $\Gamma : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is a fixed L -Lipschitz function of the indices, with d, m constants, while $p \gg n$ (and similarly for Π_0 , where use $\|\beta_i\|_0 \leq s$). Assuming $v_{DR}(y, a, z; \pi)$ is a Lipschitz function of $\pi(z)$ and since Γ is a Lipschitz function of $\langle \beta, z \rangle$, we have by a standard multi-variate Lipschitz contraction argument (and since d, m are constants), that the entropy of \mathcal{F}_Π is of the same order as the maximum entropy of each of the linear index spaces: $B_1 := \{z \rightarrow \langle \beta_i, z \rangle : \beta \in \mathbb{R}^p, \|\beta_i\|_1 \leq s\}$. Moreover, by known covering arguments (see e.g. [25], Theorem 3) that if $\|z\|_\infty \leq 1$, then:

$$\mathcal{H}_2(\epsilon, B, n) = O\left(\frac{s^2 \log(d)}{\epsilon^2}\right). \text{ Thus we get } \kappa(r, \mathcal{F}_\Pi) = O\left(s \log(n) \sqrt{\frac{\log(d)}{n}} + \frac{r}{n}\right), \text{ which}$$

leads to regret $O\left(s \log(n) \sqrt{\frac{\log(d)}{n}} + \sqrt{\frac{V_2^0 \log(1/\delta)}{n}} + h_{n,\delta}^2\right)$. In this setting, we observe that

the policy space is too large for the variance to drive the asymptotic regret. There is a leading term that remains even if the worse-case variance of policies in a small-regret slice is 0. Intuitively this stems from the high-dimensionality of the linear indices, which introduces an extra dimension of error, namely bias due to regularization. On the contrary, for exactly sparse policies $B_0 := \{z \rightarrow \langle \beta_i, z \rangle : \beta \in \mathbb{R}^p, \|\beta_i\|_0 \leq s\}$, we have that since for any possible support the entropy at scale ϵ is at most $O(s \log(1/\epsilon))$, we can take a union over all $\binom{p}{s} \leq \left(\frac{ep}{s}\right)^s$ possible sparse supports, which implies $\mathcal{H}_2(\epsilon, \mathcal{F}_\Pi, n) = O(s(\log(d/s) + \log(1/\epsilon)))$.

Thus $\kappa(r, \mathcal{F}_\Pi) = O\left(r\sqrt{\log(1/r)}\sqrt{\frac{s \log(d/s)}{n}}\right)$, leading to policy regret similar to the VC classes:

$$O\left(\sqrt{V_2^0 \log(1/V_2^0)}\sqrt{\frac{s \log(d/s)}{n}} + \sqrt{\frac{V_2^0 \log(1/\delta)}{n}} + h_{n,\delta}^2\right).$$

Remark 4 (Estimating the First Stages). *Bounds on first stage errors as a function of sample complexity measures for the first stage hypotheses spaces can be obtained by standard results on the MSE achievable by regression problems (see e.g. [18, 23]). Essentially these are bounds for the regression estimates $\hat{\theta}(z)$ and $\hat{\Sigma}(z)$, as a function of the complexity of their assumed hypothesis spaces. Since the latter is a standard statistical learning problem that is orthogonal to our main contribution, we omit technical details. Since the square loss is a strongly convex objective the rates achievable for these problems are typically fast rates on the MSE (e.g. $h_{n,\delta}^2$ is of the order $1/n$ for the*

case of parametric hypothesis spaces, and typically $o(1/\sqrt{n})$ for reproducing kernel Hilbert spaces with fast eigendecay (see e.g. [23]). Thus the term $h_{n,\delta}^2$ is of lower order. For instance, the required rates for the term $h_{n,\delta}^2$ to be of second order in our regret bounds are achievable if these nuisance regressions are ℓ_1 -penalized linear regressions and several regularity assumptions are satisfied by the data distribution, even when the dimension p of z is growing with n .

Extension: Semi-Bandit Feedback Suppose that our value function takes the form: $V(a, z) = \phi(a, z)^T \Theta_0(z) \phi(a, z)$, where $\Theta_0(z)$ is a $p \times p$ matrix and we observe semi-bandit feedback, i.e. we observe a vector Y s.t.: $\mathbb{E}[Y | a, z] = \Theta_0(z)^T \phi(a, z)$. Then we can apply our DR approach to each coordinate of Y separately.

$$V_{DR}(\pi) = \mathbb{E}_n \left[\phi(\pi(z), z)^T \left(\hat{\Theta}(z) + \hat{\Sigma}(z)^{-1} \phi(a, z) (Y^T - \phi(a, z)^T \hat{\Theta}(z)) \right) \phi(\pi(z), z) \right]$$

All the theorems in this section extend to this case, which will prove useful in our pricing application where a is the price of a set of products and Y is the vector of observed demands for each product.

4 Application: Personalized Pricing

Consider the personalized pricing of a single product. The objective is the revenue:

$$V(p, z) = p(a(z) - b(z)p)$$

where $b(z) \geq \gamma > 0$ and $a(z) + b(z)p$ gives the unknown, context-specific demand function. We assume that we observe an unbiased estimate d of demand:

$$\mathbb{E}[d | z, p] = a(z) - b(z)p$$

We want to optimize over a space of personalized pricing policies Π . If, for instance, the observational policy was homoskedastic (i.e. the exploration component was independent of the context z), we show in Appendix G that doubly robust estimators for $a(z)$ and $b(z)$ are

$$\begin{aligned} a_{DR}(z) &= \hat{a}(z) + \left(1 + \hat{g}(z) \frac{\hat{g}(z) - p}{\hat{\sigma}^2} \right) (d - \hat{a}(z) - \hat{b}(z)p) \\ b_{DR}(z) &= \hat{b}(z) + \frac{p - \hat{g}(z)}{\hat{\sigma}^2} (d - \hat{a}(z) - \hat{b}(z)p) \end{aligned}$$

where $g(z) = \mathbb{E}[p | z]$ and the variance σ^2 . Thus, in this example, we only need to estimate the mean treatment policy $\mathbb{E}[p | z]$ and the variance σ^2 .

Experimental evaluation. We empirically evaluate our framework on the personalized pricing application with synthetic data. In particular, we use simulations to assess our estimator's ability to evaluate and optimize personalized pricing functions. To do this, we compare the performance of our doubly robust estimator with (1) Direct estimator, $\langle \hat{\theta}(z), \phi(a, z) \rangle$, (2) Inverse propensity score estimator², (3) Oracle orthogonal estimator, $v_{DR}^o(x, \pi)$.

Data Generating Process. Our simulation design considers a sparse model. We assume that there are k continuous context variables distributed uniformly $z_i \sim U(1, 2)$ for $i = 1, \dots, k$ but only l of them affects demand. Let $\bar{z} = 1/l(z_1 + \dots + z_l)$. Price p and demand d are generated as $x \sim \mathcal{N}(\bar{z}, 1)$, $d = a(\bar{z}) - b(\bar{z})x + \epsilon$ and $\epsilon \sim \mathcal{N}(0, 1)$. We consider four functional forms for the demand model: (i) (Quadratic) $a(z) = 2z^2$, $b(z) = 0.6z$, (ii) (Step) $a(z) = 5\{z < 1.5\} + 6\{z > 1.5\}$, $b(z) = 0.7\{z < 1.5\} + 1.2\{z > 1.5\}$, (iii) (Sigmoid) $a(z) = 1/(1 + \exp(z)) + 3$, $b(z) = 2/(1 + \exp(z)) + 0.1$, (iv) (Linear) $a(z) = 6z$, $b(z) = z$

These functions and the data generating process ensure that the conditional expectation function of demand given z is non-negative for all z , the observed prices are positive with high probability, and the optimal prices are in the support of the observed prices. In each experiment, we generate 1000, 2000, 5000, and 10000 data points, and report results over 100 simulations. We estimate the nuisance functions using 5-fold cross-validated lasso model with polynomials of degrees up to 3 and all the two-way interactions of context variables. We present the results for two regimes: (i) Low dimensional with $k = 2, l = 1$, (ii) High dimensional with $k = 10, l = 3$.

² $\theta_{IPS}(y, a, z) = \hat{\Sigma}(z)^{-1} \phi(a, z) y$

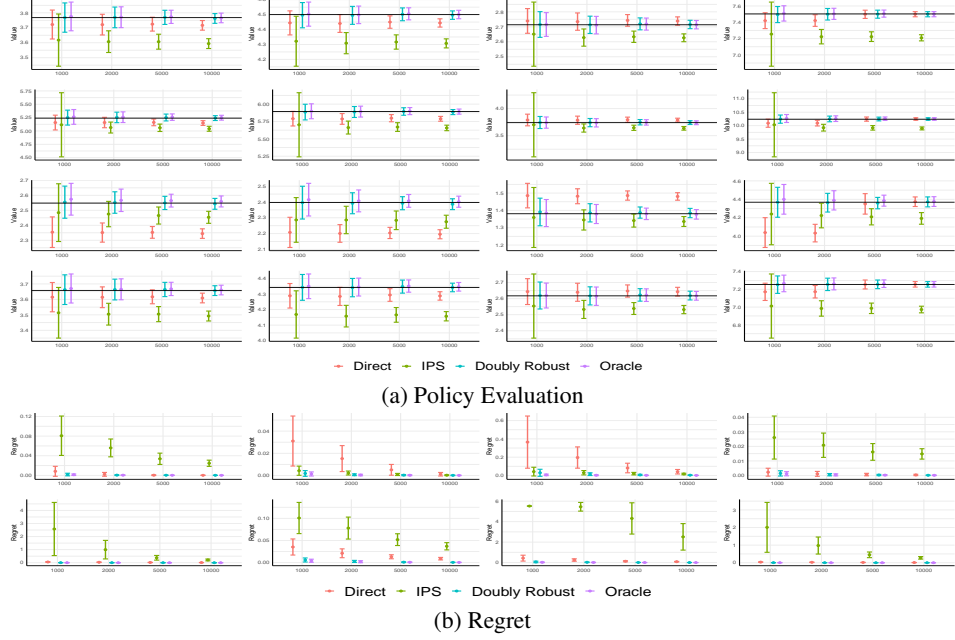


Figure 1: (a) Black line shows the true value of the policy and each line shows the mean and standard deviation of the value of the corresponding policy over 100 simulations. (b) Each line shows the mean and standard deviation of regret over 100 simulations. The top half reports the regret for a constant policy, the bottom half reports regret for a linear policy.

Policy Evaluation. For policy evaluation we consider four pricing functions: (i) Constant, $\pi(z) = 1$, (ii) Linear, $\pi(z) = z$, (iii) Threshold, $\pi(z) = 1 + 1\{z > 1.5\}$, (iv) Sin, $\pi(z) = \sin(z)$. The results for the low dimensional regime are summarized in Figure 1(a), where each row and column corresponds to a different demand function and a policy function, respectively³. The results show that, as expected, our the performance of our method is very similar to the oracle estimator and achieves a significantly better performance than the direct and inverse propensity score methods, which suffer from large biases. These results also support our claim that the asymptotic variance of the doubly robust estimate is the same as the variance of the oracle method. It is also important to point out that we obtain similar performances across two different regimes.

Regret. To investigate the regret performance of our method, we consider a constant pricing function, $\pi(z) = \gamma$ and a linear policy $\pi(z) = (\gamma_1 z_1 + \dots + \gamma_k z_k)$. We compute the optimal pricing functions in these two function spaces and report the distribution of regret in Figure 4(b) under the low dimensional regime and in Appendix H under the high dimensional regime. Across the four demand functions and two pricing functions we consider, our method achieves small regrets, comparable to the oracle method. The direct and inverse propensity methods, depending on the demand function, yield large regrets.

4.1 Quadratic Model

Finally, we consider the same simulation exercise under the assumption that an unbiased estimate of revenue rather than demand is observed. Since revenue depends on the p^2 the model is now quadratic

$$r = a(z)p - b(z)p^2 + \epsilon$$

For the data generating process we use the same functions $a(z)$ and $b(z)$ as in the personalized pricing example⁴. Figures 2 and 5 in Appendix H summarize results for policy evaluation and optimization. The overall performance of our doubly robust estimator is similar to the demand model,

³The results are very similar for the high dimensional model which are reported in Figure 4(a) in the appendix.

⁴We provide the calculation of the doubly robust estimators for this example in Appendix G.

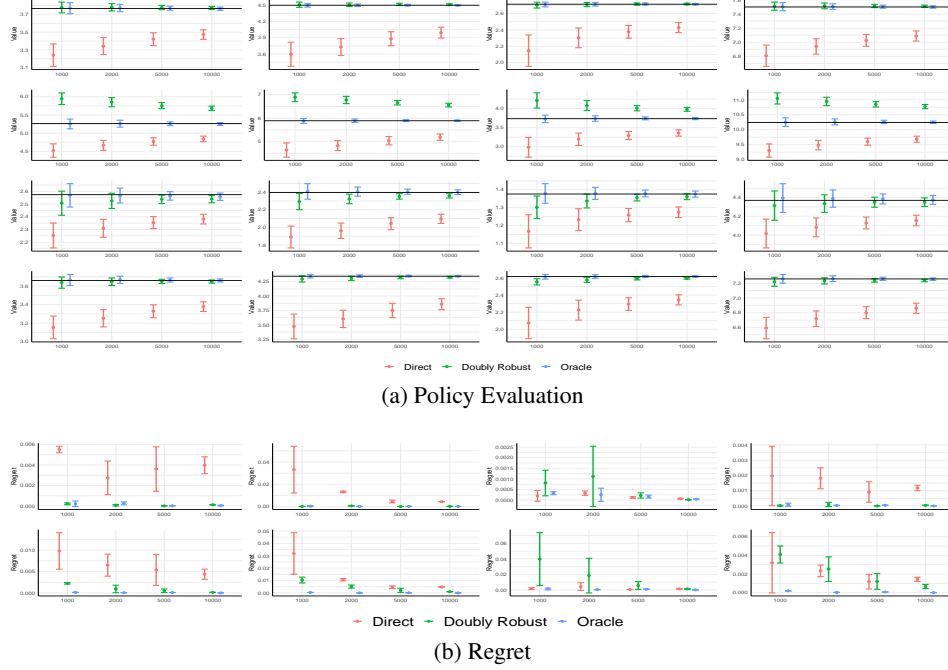


Figure 2: Quadratic, Low Dimensional Regime: (a) Black line shows the true value of the policy, each line shows the mean and standard deviation of the value of the corresponding policy over 100 simulations. (b) Each line shows the mean and standard deviation of regret over 100 simulations. We omit the results for the inverse propensity score method since they are too large to report together with the other estimates in the high dimensional regime.

and it performs better the direct model. One important difference to note is that when the sample size is small, we observe some finite sample biases for some function classes.

5 Application: Costly Resource Allocation

Motivated by a resource allocation scenario, we also analyze experimentally the special case where $\phi(a, z) = a$. Consider the case where we have p possible tasks to invest in, and we have investment costs. Each task yields a return on investment that is a linear function of the investment, but an unknown function $\theta(z)$ of the context z . Moreover, to maintain an investment portfolio of $\pi(z)$ we need to pay a known cost $C(\pi(z))$. Given a policy space $\Pi : \mathcal{Z} \rightarrow \mathbb{R}^p$, our goal is to optimize:

$$\sup_{\pi \in \Pi} \mathbb{E} [\langle \theta(z), \pi(z) \rangle - C(\pi(z))] \quad (20)$$

This falls into our framework, if we treat the offset part as of the form $\langle \theta_0(z), C(\pi(z)) \rangle$ but with a known $\theta_0(z) = 1$. So in that case we simply consider $\theta_{DR,0}(z) = \theta_0(z) = 1$. Then applying our framework we optimize:

$$\sup_{\pi \in \Pi} \mathbb{E}_n [\langle \theta_{DR}(z), \pi(z) \rangle - C(\pi(z))] \quad (21)$$

In the case of quadratic costs $C(\pi(z)) = \frac{\lambda}{2} \|\pi(z)\|_2^2$, then this boils down to exactly optimizing a square loss objective, since:

$$\inf_A \mathbb{E}_n [\|\theta_{DR}(z)/\lambda - \pi(z)\|^2] \Leftrightarrow \sup_A \mathbb{E} [\langle \theta_{DR}(z), \pi(z) \rangle] - \frac{\lambda}{2} \mathbb{E}_n [\|\pi(z)\|_2^2] \quad (22)$$

Thus policy optimization reduces to a multi-task regression problem where we are trying to predict $\theta_{DR}(z)/\lambda$ from z .⁵

⁵The above reasoning extends to heterogeneous costs across tasks e.g. $C(\pi(z)) = \sum_i c_i \pi_i(z)^2$. In this case the label of the i -th task of the multi-task regression problem is $\theta_{DR,i}(z)/c_i$ and we need to perform a weighted multi-task regression where the weight on the square loss for task i is equal c_i .

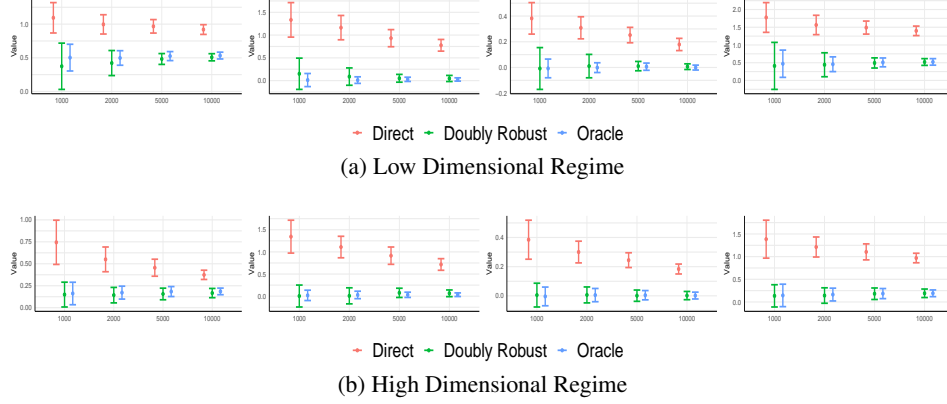


Figure 3: Costly Resource Allocation: Each line shows the mean and standard deviation of regret over 100 simulations.

We can consider sparse linear policies:

$$\Pi = \{z \rightarrow Az : \|A\|_{11} := \sum_i \|\alpha_i\|_1 \leq s\} \quad (23)$$

where α_i corresponds to the i -th row of matrix A . In this case our problem reduces to the MultiTask Lasso problem where the label is $\theta(z)/\lambda$.

Experimental Evaluation. For experimental evaluation we consider a model with two tasks, a_1 and a_2 :

$$y = a(z)a_1 + b(z)a_2 + \epsilon$$

We use the same distributions and functions, $a(z)$ and $b(z)$, given above for the pricing application. To estimate the optimal allocation and its regret, we run a 5-fold cross validated MultiTask Lasso algorithm and set $\lambda = 1$. We report the distribution of return on investment obtained from different models in Figure (3). The results suggest that doubly robust method achieves a significantly lower regret than the direct method in both regimes and its performance is similar to the oracle method ⁶.

References

- [1] Susan Athey and Stefan Wager. Efficient policy learning. *arXiv preprint arXiv:1702.02896*, 2017.
- [2] Alina Beygelzimer and John Langford. The offset tree for learning with partial labels. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 129–138. ACM, 2009.
- [3] Peter J Bickel, Chris AJ Klaassen, Peter J Bickel, Y Ritov, J Klaassen, Jon A Wellner, and YA’Acov Ritov. *Efficient and adaptive estimation for semiparametric models*, volume 4. Johns Hopkins University Press Baltimore, 1993.
- [4] Gary Chamberlain. Efficiency bounds for semiparametric regression. *Econometrica: Journal of the Econometric Society*, pages 567–596, 1992.
- [5] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.
- [6] Victor Chernozhukov, Juan Carlos Escanciano, Hidehiko Ichimura, Whitney K Newey, and James M Robins. Locally robust semiparametric estimation. *arXiv preprint arXiv:1608.00033*, 2016.

⁶For comparison, the value achieved by best-in-class policy is 22.2 in low dimensional regime and ? in high dimensional regime. We omit the inverse propensity score regrets since they are too large to report together with other estimates

- [7] Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on Machine Learning*, pages 1097–1104. Omnipress, 2011.
- [8] Robert F Engle, Clive WJ Granger, John Rice, and Andrew Weiss. Semiparametric estimates of the relation between weather and electricity sales. *Journal of the American statistical Association*, 81(394):310–320, 1986.
- [9] Dylan J Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*, 2019.
- [10] Bryan S Graham and Cristine Campos de Xavier Pinto. Semiparametrically efficient estimation of the average linear regression function. Working paper, National Bureau of Economic Research, 2018.
- [11] Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- [12] Nathan Kallus and Angela Zhou. Policy evaluation and optimization with continuous treatments. *arXiv preprint arXiv:1802.06037*, 2018.
- [13] Toru Kitagawa and Aleksey Tetenov. Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica*, 86(2):591–616, 2018.
- [14] Akshay Krishnamurthy, John Langford, Aleksandrs Slivkins, and Chicheng Zhang. Contextual bandits with continuous actions: Smoothing, zooming, and adapting. *arXiv preprint arXiv:1902.01520*, 2019.
- [15] Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample variance penalization. In *The 22nd Conference on Learning Theory (COLT)*, 2009.
- [16] Whitney K Newey. Semiparametric efficiency bounds. *Journal of applied econometrics*, 5(2):99–135, 1990.
- [17] Min Qian and Susan A Murphy. Performance guarantees for individualized treatment rules. *Annals of statistics*, 39(2):1180, 2011.
- [18] Alexander Rakhlin, Karthik Sridharan, Alexandre B Tsybakov, et al. Empirical entropy, minimax regret and minimax risk. *Bernoulli*, 23(2):789–824, 2017.
- [19] Peter M Robinson. Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, pages 931–954, 1988.
- [20] Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning*, pages 814–823, 2015.
- [21] Mark J Van der Laan and Sherri Rose. *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media, 2011.
- [22] A. W. Van Der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Series, March 1996.
- [23] Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.
- [24] Jeffrey M Wooldridge. Estimating average partial effects under conditional moment independence assumptions. Working paper, cemmap working paper, 2004.
- [25] Tong Zhang. Covering number bounds of certain regularized linear function classes. *Journal of Machine Learning Research*, 2(Mar):527–550, 2002.
- [26] Yingqi Zhao, Donglin Zeng, A John Rush, and Michael R Kosorok. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012.
- [27] Xin Zhou, Nicole Mayer-Hamblett, Umer Khan, and Michael R Kosorok. Residual weighted learning for estimating individualized treatment rules. *Journal of the American Statistical Association*, 112(517):169–187, 2017.
- [28] Zhengyuan Zhou, Susan Athey, and Stefan Wager. Offline multi-action policy learning: Generalization and optimization. *arXiv preprint arXiv:1810.04778*, 2018.

A Proof of Universal Orthogonality Lemma

We first start by defining a sufficient condition for the notion of *universal orthogonality* of a loss function, as defined by [9]. A loss function $L(\pi; h) = \mathbb{E}[\ell(x, \pi(z); h(z))]$ is universally orthogonal with respect to h if for any $\pi \in \Pi$:

$$\mathbb{E}[\nabla_{h(z), \pi(z)} \ell(x, \pi(z); h_0(z)) \mid z] = 0 \quad (24)$$

where h_0 is the true value of the nuisance parameter h .

Lemma 3. *The loss function $L(\pi; h) = -\mathbb{E}[\langle \theta_{DR}(y, a, z), \phi(\pi(z), z) \rangle]$ is universally orthogonal with respect to $h = (\theta, \Sigma)$.*

Proof. We show that the population loss function that corresponds to the doubly robust estimate, satisfies the universal orthogonality condition. For simplicity of notation let $K(z) = \Sigma(z)^{-1}$. Then the population loss is:

$$V_{DR}^0(\pi; \hat{\theta}, \Sigma^{-1}) = \mathbb{E} \left[\left\langle \hat{\theta}(z) + \Sigma^{-1}(z) \phi(a, z) (y - \langle \hat{\theta}(z), \phi(a, z) \rangle), \phi(\pi(z), z) \right\rangle \right]$$

Let:

$$\beta(a, z, \xi, K) = \xi + K \phi(a, z) (y - \langle \xi, \phi(a, z) \rangle)$$

Observe that:

$$V_{DR}^0(\pi; \hat{\theta}, \Sigma^{-1}) = \mathbb{E} \left[\left\langle \beta(a, \hat{\theta}(z), \Sigma^{-1}(z)), \phi(\pi(z), z) \right\rangle \right]$$

To show universal orthogonality it suffices to show that:

$$\mathbb{E} [\nabla_{\xi, K} \beta(a, z, \theta_0(z), \Sigma_0^{-1}(z)) \mid z] = 0$$

This follows easily by simple algebraic manipulations:

$$\begin{aligned} \mathbb{E} [\nabla_{\xi} \beta(a, z, \theta_0(z), \Sigma_0^{-1}(z)) \mid z] &= \mathbb{E} [\mathbb{I} - \Sigma_0^{-1}(z) \phi(a, z) \phi(a, z)^T \mid z] \\ &= \mathbb{I} - \Sigma_0^{-1}(z) \mathbb{E} [\phi(a, z) \phi(a, z)^T \mid z] = \mathbb{I} - \Sigma_0^{-1}(z) \Sigma_0(z) = 0 \end{aligned}$$

and

$$\mathbb{E} [\nabla_{K_{ij}} \beta(a, \theta_0(z), \Sigma_0^{-1}(z)) \mid z] = \mathbb{E} [\phi_j(a, z) (y - \langle \theta_0(z), \phi(a, z) \rangle) \mid z]$$

Now observe that since $\theta_0(z)$ is the minimizer of the conditional squared loss, taking the first order condition implies:

$$\begin{aligned} \mathbb{E}[(V_0(a, z) - \langle \theta_0(z), \phi(a, z) \rangle) \phi(a, z) \mid z] &= 0 \iff \\ \mathbb{E}[V_0(a, z) \phi(a, z) \mid z] &= \mathbb{E}[\langle \theta_0(z), \phi(a, z) \rangle \phi(a, z) \mid z] \end{aligned}$$

Moreover:

$$\mathbb{E}[y \phi(a, z) \mid z] = \mathbb{E}[\mathbb{E}[y \mid a, z] \phi(a, z)] = \mathbb{E}[V_0(a, z) \phi(a, z)]$$

Combining the two yields:

$$\mathbb{E} [\phi(a, z) (y - \langle \theta_0(z), \phi(a, z) \rangle) \mid z] = 0$$

which implies orthogonality with respect to K . \square

B Proof of Main Regret Theorem 1

We first consider an arbitrary empirical loss minimization problem of the form:

$$f_n = \arg \min_{f \in \mathcal{F}} \mathbb{E}_n[f(x)] := \frac{1}{n} \sum_{i=1}^n f(x_i) \quad (25)$$

where $x_i \in \mathcal{X}$ are i.i.d. drawn from an unknown distribution and \mathcal{X} is an arbitrary data space. Throughout the section we will assume that: $\sup_{f \in \mathcal{F}} |f(x)| \leq 1$. All the results can be generalized

to the case of $\sup_{f \in \mathcal{F}} |f(x)| \leq R$, for some arbitrary R , by simply first re-scaling the losses, and then invoking the theorems of this section.

We will also make the following preliminary definitions. For any function f we denote with $\|f\|_2 = \sqrt{\mathbb{E}[f(x)^2]}$, the standard L^2 norm and with $\|f\|_{2,n} = \sqrt{\mathbb{E}_n[f(x)^2]}$ its empirical analogue. The localized Rademacher complexity is the defined as:

$$\mathcal{R}(r, \mathcal{F}) = \mathbb{E}_{\epsilon, x_{1:n}} \left[\sup_{f \in \mathcal{F}: \|f\|_2 \leq r} \frac{1}{n} \sum_{i=1}^n \epsilon_i f(x_i) \right] \quad (26)$$

where ϵ_i are independent Rademacher variables that take values $\{-1, 1\}$ with equal probability.

Furthermore, we define the empirical entropy of a function class $H_2(\epsilon, \mathcal{F}, n)$ as the largest value, over the choice of n samples, of the logarithm of the size of the smallest empirical ϵ -cover of \mathcal{F} on the samples with respect to the $\|\cdot\|_{2,n}$ norm. Finally, we consider the empirical entropy integral defined as:

$$\kappa(r, \mathcal{F}) = \inf_{\alpha \geq 0} \left\{ 4\alpha + 10 \int_{\alpha}^r \sqrt{\frac{\mathcal{H}_2(\epsilon, \mathcal{F}, n)}{n}} d\epsilon \right\}, \quad (27)$$

Throughout this section we will make the following benign assumption that essentially makes the problem *learnable*:

ASSUMPTION 1. *The function class satisfies that for any constant r , $\kappa(r, \mathcal{F}) \rightarrow 0$ as $n \rightarrow \infty$*

We will use the following theorems from the prior work of [9] as a starting point as they are formalized in manner convenient for our problem.

Theorem 4 (Foster, Syrgkanis [9], Theorem 4). *Consider any function class $\mathcal{F} : \mathcal{X} \rightarrow [-1, 1]$ and let f_n be the outcome of the constrained ERM. Pick any $f_* \in \mathcal{F}$ and let $r = \sup_{f \in \mathcal{F}} \|f - f_*\|_2$. Then for some constants C_1, C_2 and for any $\delta > 0$, w.p. $1 - \delta$:*

$$\begin{aligned} \mathbb{E}[f_n(x) - f_*(x)] &\leq C_1 \left(\mathcal{R}(r, \mathcal{F} - f_*) + r \sqrt{\frac{\log(1/\delta)}{n}} + \frac{\log(1/\delta)}{n} \right) \\ &\leq C_1 C_2 \left(\kappa(r, \mathcal{F}) + r \sqrt{\frac{\log(1/\delta)}{n}} + \frac{\mathcal{H}_2(r, \mathcal{F}, n)}{n} + \frac{\log(1/\delta)}{n} \right). \end{aligned}$$

Lemma 5 (Foster, Syrgkanis [9], Lemma 4). *Consider a function class $\mathcal{F} : \mathcal{X} \rightarrow [-1, 1]$ and pick any $f_* : \mathcal{X} \rightarrow [-1, 1]$ (not necessarily in \mathcal{F}). Moreover, let:*

$$Z_n(r) = \sup_{f \in \mathcal{F}: \|f - f_*\|_2 \leq r} |\mathbb{E}_n[f(x) - f_*(x)] - \mathbb{E}[f(x) - f_*(x)]| \quad (28)$$

Then for some constant C_3 and for any $\delta > 0$, w.p. $1 - \delta$:

$$Z_n(r) \leq C_3 \left(\mathcal{R}(r, \mathcal{F} - f_*) + r \sqrt{\frac{\log(1/\delta)}{n}} + \frac{\log(1/\delta)}{n} \right)$$

Our goal is to replace r in the latter Theorem with the worst-case variance of the functions $f \in \mathcal{F}$ in a small “regret”-ball around the optimal. We will achieve this by considering a slight modification of the ERM algorithm. In particular, we will split the data in half, and we will use one half as a *regularization sample* and the other half as the *training sample*. In particular, we will find the optimal function on the training sample, within the class of functions that also have relatively small regret on the regularization sample.

Out-of-Sample Regularized ERM Consider the following algorithm:

- We split the samples S in two parts S_1, S_2 and let $\mathbb{E}_{n_1}[\cdot]$ and $\mathbb{E}_{n_2}[\cdot]$ denote the corresponding empirical expectations.
- We run ERM over \mathcal{F} on the first half and let f_1 be the outcome.
- Then we define the class of functions that have the constraint that they don’t achieve much worse value than f_1 on the first half, i.e. we regularize policies based on their regret on the first half. More formally, for some constant μ_n to be defined later:

$$\mathcal{F}_2 = \{f \in \mathcal{F} : \mathbb{E}_{n_1}[f(x) - f_1(x)] \leq \mu_n\} \quad (29)$$

- Then we run constrained ERM on the second sample over the function space \mathcal{F}_2 :

$$f_2 = \arg \min_{f \in \mathcal{F}_2} \mathbb{E}_{n_2}[f(x)] \quad (30)$$

Theorem 6 (Variance-Based Regret). *Let $f_* = \arg \min_{f \in \mathcal{F}} \mathbb{E}[f(x)]$, $r = \sup_{f \in \mathcal{F}} \|f\|_2$ and choose $\mu_n = C \left(\kappa(r, \mathcal{F}) + r \sqrt{\frac{\log(6/\delta)}{n}} + \frac{\mathcal{H}_2(r, \mathcal{F}, n)}{n} + \frac{\log(6/\delta)}{n} \right)$, with $C = 8 \max\{C_1 C_2, C_3 C_2\}$. Then, w.p. $1 - \delta$ over the sample S , the outcome f_2 of the Out-of-Sample Regularized ERM satisfies:*

$$\mathbb{E}[f_2(x) - f_*(x)] = O \left(\kappa(\sqrt{V_2}, \mathcal{F}_*(\mu_n)) + \sqrt{\frac{V_2 \log(3/\delta)}{n}} \right) \quad (31)$$

with: $\mathcal{F}_*(\mu_n) = \{f \in \mathcal{F} : \mathbb{E}[f(x) - f_*(x)] \leq \mu_n\}$ and $V_2 = \sup_{f \in \mathcal{F}_*(\mu_n)} \text{Var}(f(x) - f_*(x))$. Moreover, the expected regret, in expectation over the samples S_1, S_2 is also of order $O \left(\kappa(\sqrt{V_2}, \mathcal{F}) + \sqrt{\frac{V_2}{n}} \right)$.

Proof. First we argue that w.p. $1 - \delta/6$, $f_* \in \mathcal{F}_2$. By the choice of μ_n and Theorem 4, we know that w.p. $1 - \delta/4$ over the randomness of sample S_1 :

$$\mathbb{E}[f_1(x) - f_*(x)] \leq \mu_n/2 \quad (32)$$

Moreover, by Lemma 5, w.p. $1 - \delta/6$ over the randomness of sample S_1 :

$$\sup_{f \in \mathcal{F}} |\mathbb{E}_{n_1}[f(x) - f_*(x)] - \mathbb{E}[f(x) - f_*(x)]| \leq \mu_n/2$$

Combining the latter two properties we have, w.p. $1 - \delta/3$:

$$|\mathbb{E}_{n_1}[f_*(x) - f_1(x)]| \leq |\mathbb{E}[f_*(x) - f_1(x)]| + \mu_n/2 \leq \mu_n$$

Thus in this event, $f_* \in \mathcal{F}_2$.

Applying Theorem 4 for the last stage of the algorithm with function space \mathcal{F}_2 and conditioning on the event that the first stage sample is such that $f_* \in \mathcal{F}_2$, we have that with probability $1 - \delta/3$ over the randomness of the second sample:

$$\mathbb{E}[f_2(x) - f_*(x)] = C_1 C_2 \left(\kappa(r_2, \mathcal{F}_2) + r_2 \sqrt{\frac{\log(3/\delta)}{n}} + \frac{\mathcal{H}_2(r_2, \mathcal{F}_2, n)}{n} + \frac{\log(3/\delta)}{n} \right)$$

where $r_2 = \sup_{f \in \mathcal{F}_2} \|f\|_2$. Thus by a union bound we get that with probability $1 - 2\delta/3$ over the randomness of both samples, the latter bound holds.

Observe that for $f \in \mathcal{F}_2$, by Lemma 5, w.p. $1 - \delta/6$ over the first sample:

$$\sup_{f \in \mathcal{F}} |\mathbb{E}_{n_1}[f(x) - f_1(x)] - \mathbb{E}[f(x) - f_1(x)]| \leq 2 \sup_{f \in \mathcal{F}} |\mathbb{E}_{n_1}[f(x)] - \mathbb{E}[f(x)]| \leq \mu_n/2$$

Thus w.p. $1 - \delta/6$, \mathcal{F}_2 is a subset of the class:

$$\mathcal{F}_2^0 = \{f \in \mathcal{F} : |\mathbb{E}[f(x) - f_1(x)]| \leq \mu_n/2\} \quad (33)$$

Moreover, since f_1 has small regret, we know by the triangle inequality, for all $f \in \mathcal{F}_2^0$, w.p. $1 - \delta/3$:

$$|\mathbb{E}[f(x) - f_*(x)]| \leq |\mathbb{E}[f(x) - f_1(x)]| + |\mathbb{E}[f_1(x) - f_*(x)]| \leq \mu_n$$

Thus w.p. $1 - \delta/3$, \mathcal{F}_2^0 is in turn a subset of the function space:

$$\mathcal{F}_*(\mu_n) = \{f \in \mathcal{F} : |\mathbb{E}[f(x) - f_*(x)]| \leq \mu_n\}$$

which is a space of policies with regret at most μ_n .

Thus we have that w.p. $1 - \delta/3$ over the first sample:

$$\begin{aligned} r_2^2 &= \sup_{f \in \mathcal{F}_2} \mathbb{E}[(f(x) - f_*(x))^2] \leq \sup_{f \in \mathcal{F}_*(\mu_n)} \mathbb{E}[(f(x) - f_*(x))^2] \\ &= \sup_{f \in \mathcal{F}_*(\mu_n)} (\text{Var}(f(x) - f_*(x)) + \mathbb{E}[f(x) - f_*(x)]^2) \\ &\leq \sup_{f \in \mathcal{F}_*(\mu_n)} \text{Var}(f(x) - f_*(x)) + \mu_n^2 \end{aligned}$$

We thus have that:

$$r_2 = \sqrt{\sup_{f \in \mathcal{F}_*(\mu_n)} \text{Var}(f(x) - f_*(x))} + 2\mu_n = \sqrt{V_2} + 2\mu_n \quad (34)$$

Combining the latter with the regret bound for f_2 (excluding lower order terms in n) we have that w.p. $1 - \delta$:

$$\mathbb{E}[f_2(x) - f_*(x)] = O \left(\kappa(\sqrt{V_2} + 2\mu_n, \mathcal{F}_*(\mu_n)) + \sqrt{\frac{V_2 \log(3/\delta)}{n}} \right)$$

Moreover, using the concavity of the entropy integral with respect to its first argument, we have that:

$$\kappa(\sqrt{V_2} + 2\mu_n, \mathcal{F}) \leq \kappa(\sqrt{V_2}, \mathcal{F}_*(\mu_n)) + 2\mu_n \sqrt{\frac{H_2(\sqrt{V_2}, \mathcal{F}, n)}{n}} \quad (35)$$

Since $\kappa(r, \mathcal{F}) \rightarrow 0$, we have that $\mu_n = o(1)$ and $H_2(\sqrt{V_2}, \mathcal{F}, n)$ is a constant. Thus, the second term decays faster than $1/\sqrt{n}$ and hence is asymptotically negligible. Thus we get:

$$\mathbb{E}[f_2(x) - f_*(x)] = O \left(\kappa(\sqrt{V_2}, \mathcal{F}_*(\mu_n)) + \sqrt{\frac{V_2 \log(1/\delta)}{n}} \right)$$

The expected regret bound follows by standard arguments by simply integrating the above high probability bound. \square

Going back to our policy learning problem, let $x = (y, a, z)$ and:

$$v_{DR}(x; \pi) = \langle \theta_{DR}(y, a, z), \phi(\pi(z), z) \rangle \quad (36)$$

be the doubly robust proxy value at every sample x and policy π . Then we can apply this general theorem to the policy learning problem where, $x = (y, a, z)$ and function space:

$$\mathcal{F}_\Pi = \{-v_{DR}(\cdot; \pi) : \pi \in \Pi\} \quad (37)$$

Then Theorem 6 yields the following corollary:

Corollary 7 (Variance-Based Policy Regret). *Let $\pi_* = \arg \max_{\pi \in \Pi} \mathbb{E}[v_{DR}(x; \pi)]$, $r = \sup_{\pi \in \Pi} \sqrt{\mathbb{E}[v_{DR}(z; \pi)^2]}$, $\mu_n = \Theta \left(\kappa(r, \mathcal{F}_\Pi) + r \sqrt{\frac{\log(1/\delta)}{n}} \right)$ and*

$$V_2 = \sup_{\pi \in \Pi: \mathbb{E}[v_{DR}(x; \pi_*) - v_{DR}(x; \pi)] \leq \mu_n} \text{Var}(v_{DR}(x; \pi) - v_{DR}(x; \pi_*)). \quad (38)$$

Then the policy π_2 returned by the out-of-sample regularized ERM, satisfies w.p. $1 - \delta$ over the randomness of S :

$$\mathbb{E}[v_{DR}(\pi_*) - v_{DR}(\pi_2)] = O \left(\kappa(\sqrt{V_2}, \mathcal{F}_\Pi) + \sqrt{\frac{V_2 \log(1/\delta)}{n}} \right) \quad (39)$$

and expected regret $O \left(\kappa(\sqrt{V_2}, \mathcal{F}_\Pi) + \sqrt{\frac{V_2}{n}} \right)$.

To arrive at our final theorem, we also need to account for the difference between $\mathbb{E}[v_{DR}(x; \pi)]$ and $V(\pi)$. This difference essentially stems from the error in the nuisance estimates, which introduce an error in $\theta_{DR}(y, a, z)$, such that $\mathbb{E}[\theta_{DR}(y, a, z) | z] \neq \theta(z)$. However, we can invoke the orthogonality of the doubly robust estimator and the general theorem of [9] on generalization bounds of orthogonal losses to get:

Lemma 8. *For any policy $\pi_0 \in \Pi$, let $\hat{\pi}$ be the outcome of any possibly randomized algorithm that satisfies w.p. $1 - \delta/2$ a regret bound on the doubly robust objective, i.e. $\mathbb{E}[v_{DR}(x; \pi_0) - v_{DR}(x; \hat{\pi})] \leq R_{n, \delta}$. Moreover, suppose that the nuisance estimates satisfy a mean-squared error bound*

$$\max \left\{ \mathbb{E}[(\hat{\theta}(z) - \theta_0(z))^2], \mathbb{E}[\|\hat{\Sigma}(z) - \Sigma_0(z)\|_{F_{ro}}^2] \right\} := \chi_n^2 \quad (40)$$

Then w.p. $1 - \delta$ over the randomness of the policy sample:

$$V(\pi_0) - V(\hat{\pi}) \leq O(R_{n, \delta} + \chi_n^2) \quad (41)$$

Proof. By Lemma 3 we have that the loss function $-\mathbb{E}[v_{DR}(x; \pi)]$ is universally orthogonal as defined in [9]. Moreover, the loss is smooth with respect to the outputs of the nuisance functions and hence the second order derivatives of the loss with respect to the outputs of the nuisance functions are bounded. Thus the lemma follows by Theorem 2 of [9]. \square

If we assume that the nuisance estimation algorithm guarantees that w.p. $1 - \delta$, $\chi_n^2 \leq h_{n,\delta}^2$ then observe that combining Corollary 7 and Lemma 8, we get that for any policy π_0 , the policy π_2 of the out-of-sample regularized ERM satisfies, w.p. $1 - \delta$:

$$V(\pi_0) - V(\pi_2) \leq O\left(\kappa(\sqrt{V_2}, \mathcal{F}_\Pi) + \sqrt{\frac{V_2 \log(1/\delta)}{n}} + h_{n,\delta}^2\right)$$

Similarly, if we assume that the nuisance estimation algorithm satisfies $\mathbb{E}[\chi_n^2] \leq h_n^2$, then:

$$\mathbb{E}[V(\pi_0) - V(\pi_2)] \leq O\left(\kappa(\sqrt{V_2}, \mathcal{F}_\Pi) + \sqrt{\frac{V_2 \log(1/\delta)}{n}} + h_n^2\right)$$

We continue by proving the probabilistic regret bound of the theorem and the in-expectation bound follows analogously.

Finally, we need to account for the error introduced by the nuisance errors on the quantity V_2 , so as to connect it with the semi-parametric efficiency variance of each policy, i.e.:

$$\text{Var}(v_{DR}^0(x; \pi)) \quad (42)$$

where $v_{DR}^0(x; \pi) = \langle \theta_{DR}^0(y, a, z), \phi(\pi(z), z) \rangle$, and $\theta_{DR}^0(y, a, z)$ is the analogue of the doubly robust function, $\theta_{DR}(y, a, z)$, evaluated at the true nuisance functions. Moreover, we want our the “regret slice” to be with respect to the true regret, i.e. we want to depend on the variance of policies that satisfy:

$$V(\pi_*^0) - V(\pi) := \mathbb{E}[v_{DR}^0(x; \pi_*^0) - v_{DR}^0(x; \pi)] \leq \mu'_n \quad (43)$$

where $\pi_*^0 = \arg \max_{\pi \in \Pi} V(\pi)$. We prove such a property in the following lemma:

Lemma 9. *Consider the setting of Corollary 7. Suppose that the mean squared error of the nuisance estimates is upper bounded w.p. $1 - \delta$ by $h_{n,\delta}^2$ and let $\epsilon_n = \mu_n + h_{n,\delta}^2$. Then:*

$$V_2^0 = \sup_{\pi, \pi' \in \Pi_*(\epsilon_n)} \text{Var}(v_{DR}^0(x; \pi) - v_{DR}^0(x; \pi')) \quad (44)$$

Then $V_2 \leq V_2^0 + O(h_{n,\delta})$.

Proof. First observe that by Lemma 8 with $\pi_0 = \pi_*$ and $\hat{\pi} = \pi$ (for any $\pi \in \mathcal{F}_\Pi^2$), we have that:

$$\mathbb{E}[v_{DR}(\pi_*) - v_{DR}(\pi)] \leq \mu_n \implies V(\pi_*) - V(\pi) \leq \mu_n + O(h_{n,\delta}^2)$$

Similarly if we let $\pi_*^0 = \arg \max_{\pi \in \Pi} \mathbb{E}[v_{DR}^0(x; \pi)] := V(\pi)$, then observe that by definition of π_* : $\mathbb{E}[v_{DR}(x; \pi_*^0) - v_{DR}(x; \pi_*)] \leq 0$. Thus applying again Lemma 8 with $\pi_0 = \pi_*^0$ and $\hat{\pi} = \pi_*$:

$$\mathbb{E}[v_{DR}(\pi_*^0) - v_{DR}(\pi_*)] \leq 0 \implies V(\pi_*^0) - V(\pi_*) \leq O(h_{n,\delta}^2)$$

Let $\Pi_*^0(\epsilon) = \{\pi \in \Pi : V(\pi_*^0) - V(\pi) \leq \epsilon\}$ and let $\epsilon_n = O(\mu_n + h_{n,\delta}^2)$. Thus we have that:

$$V_2 \leq \sup_{\pi \in \Pi_*(\epsilon_n)} \text{Var}(v_{DR}(x; \pi) - v_{DR}(x; \pi_*))$$

Moreover, observe that $\pi_* \in \Pi_*^0(\epsilon_n)$. Hence:

$$V_2 \leq \sup_{\pi, \pi' \in \Pi_*(\epsilon_n)} \text{Var}(v_{DR}(x; \pi) - v_{DR}(x; \pi'))$$

Moreover, by Lipschitzness of $\theta_{DR}(y, a, z)$ on the output of the nuisance functions, we also have that for any $\pi, \pi' \in \Pi(\epsilon_n)$:

$$\text{Var}(v_{DR}(x; \pi) - v_{DR}(x; \pi')) \leq \text{Var}(v_{DR}^0(x; \pi) - v_{DR}^0(x; \pi')) + O(h_{n,\delta}) \quad (45)$$

Hence, if we denote with:

$$V_2^0 = \sup_{\pi, \pi' \in \Pi_*(\epsilon_n)} \text{Var}(v_{DR}^0(x; \pi) - v_{DR}^0(x; \pi'))$$

Then we conclude that:

$$V_2 = V_2^0 + O(h_{n,\delta})$$

□

Invoking Lemma 9 and the concavity of the entropy integral function we get:

$$V(\pi_*^0) - V(\hat{\pi}) \leq O\left(\kappa(\sqrt{V_2^0}, \mathcal{F}_\Pi) + \sqrt{\frac{V_2^0 \log(1/\delta)}{n}} + h_{n,\delta}^2 + h_{n,\delta} \frac{1}{\sqrt{n}}\right) \quad (46)$$

Since $h_{n,\delta} = o(1)$, the last term is of lower order. This concludes the proof of the main regret Theorem 1.

C Review of Semi-parametric Efficiency Bounds

In this section, we review the theory of semi-parametric efficiency bounds studied in [16] and [3].

C.1 Definitions

Definition 1 (Mean Square Differentiability). *Let $f(x; \eta)$ denote the probability density function of a random variable x where $\eta \in H$ is a finite dimensional parameter. $f(x; \eta)^{1/2}$ is μ -mean square continuously differentiable with respect to η on H with derivative $f_\eta(x; \eta)$ if for each $\eta \in H$ $\int \|f_\eta(x; \eta)\|^2 d\mu$ is finite, and for every $\eta_i \rightarrow \eta$ with $\int \|f_\eta(x; \eta_i) - f_\eta(x; \eta)\|^2 d\mu \rightarrow 0$*

$$\int \left(f(x; \eta_i)^{1/2} - f(x; \eta)^{1/2} - f_\eta(x; \eta)'(\eta_i - \eta) \right)^2 d\mu / \|\eta_i - \eta\|^2 \rightarrow 0$$

Definition 2 (Smoothness). *$f(x; \eta)$ is smooth if (i) $\eta \in H$, H is open; (ii) there is a measure μ dominating $f(x; \eta)$ for $\eta \in H$ such that $f(x; \eta)$ is continuous on H a.s. μ ; (iii) $f(x; \eta)^{1/2}$ is mean square differentiable.*

Definition 3 (Score and Information Matrix). *For smooth $f(x; \eta)$ the score for η is defined as*

$$S_\eta(x; \eta) := 2 \frac{f_\eta(x; \eta)}{f(x; \eta)}$$

in the support of x and the information matrix is

$$\mathcal{I}(\eta) = \int S_\eta S_\eta' f(x; \eta) d\mu.$$

Definition 4 (Regularity). *A likelihood function $f(x; \eta)$, $\eta \in H$, is regular if it is smooth and information matrix is non-singular in H . The efficiency bound of a regular model is given by Cramer-Rao bound and equals $\mathcal{I}(\eta)^{-1}$.*

Definition 5 (Linearity). *Define a set \mathcal{T} to be linear if $as_1 + bs_2 \in \mathcal{T}$ for all real scalars a and b and elements s_1 and s_2 of \mathcal{T} .*

C.2 Derivation of the Efficiency Bound

Let data (x_1, \dots, x_n) consist of i.i.d copies of the random vector (y, a, z) . A semi-parametric model consists of a parameter vector α and a set of restrictions on the joint behavior of observables. In our model, the restrictions are given by the first order conditions of the linear projection

$$\mathbb{E}[(y - \langle \theta_0(z), \phi(a, z) \rangle) \phi(a, z) | z] = 0$$

and the parameter is

$$\alpha = \int \langle \theta(z), \phi(\pi(z), z) \rangle f(z) dz$$

where $f(z)$ denotes the probability distribution function of z . First, we provide the definition of a parametric submodel.

Definition 6 (Parametric Submodel). *For estimators with i.i.d data, a parametric submodel corresponds to a parameter vector η and a likelihood function $\ell(x|\eta)$ for a single observation that satisfies the semi-parametric restrictions.*

A parametric submodel is a subset of the model distributions satisfying the semi-parametric assumptions. The reason parametric submodels are useful in analyzing semi-parametric efficiency is that for parametric models, the Cramer-Rao bound gives the lower bound on the variance of estimators of a parameter under some regulatory conditions. Since semi-parametric models impose weaker restrictions than any parametric model, it is natural to expect that the asymptotic variance of a semi-parametric model is no smaller than the bound for the parametric model.

In a parametric submodel, our parameter of interest can be written as

$$\alpha = \int \langle \theta(z; \eta), \phi(\pi(z), z) \rangle f(z; \eta) dz \quad (47)$$

Next, we define the semi-parametric efficient bounds.

Definition 7 (Semi-parametric Efficiency Bound). *The semi-parametric efficiency bound of a semi-parametric estimator is defined as the supremum of the Cramer-Rao bounds for all regular parametric submodels.*

This definition is intuitive because any semi-parametric estimator that is consistent and asymptotically normal cannot have a lower variance than the supremum of Cramer-Rao bounds. The regulatory conditions defined in Section C.1 guarantee that the Cramer-Rao bound is well-defined and gives an asymptotic efficiency bound.

To be able to obtain the Cramer-Rao bound for the parameter of interest under a parametric submodel, the parameter must be pathwise differentiable.

Definition 8 (Pathwise Differentiability). *A parameter α is pathwise-differentiable if $\alpha(\eta)$ is differentiable for all smooth parametric submodels and there exists $q \times 1$ random vector d such that $\mathbb{E}[d'd]$ is finite and for all regular parametric submodels*

$$\frac{\partial \alpha(\eta_0)}{\partial \eta} = \mathbb{E}[dS'_\eta]$$

where η_0 denotes the true value of the parameter in the sense that $\ell(x|\eta_0)$ corresponds to the likelihood function that generates the data.

Pathwise differentiability of a parameter is a weak condition because, by Riesz representation theorem, a parameter is pathwise-differentiable if it can be written as a functional that is mean-square continuous. From the definition of α in Equation (47) it is easy to see that α is pathwise-differentiable by Riesz representation theorem.

For a pathwise-differentiable parameter, the Cramer-Rao bound can be written as a function of the pathwise derivative using the Delta method.

$$\begin{aligned} \text{Var}(\alpha(\eta_0)) &= \frac{\partial \alpha(\eta_0)}{\partial \eta} (\mathbb{E}[S_\eta S'_\eta])^{-1} \frac{\partial \alpha(\eta_0)'}{\partial \eta} \\ &= \mathbb{E}[dS'_\eta] (\mathbb{E}[S_\eta S'_\eta])^{-1} \mathbb{E}[S_\eta d'] \end{aligned}$$

We can write $\text{Var}(\alpha(\eta))$ as a second moment of a random variable as follows

$$\begin{aligned} \text{Var}(\alpha(\eta_0)) &= \mathbb{E}[dS'_\eta] (\mathbb{E}[S_\eta S'_\eta])^{-1} \mathbb{E}[S_\eta d'] \\ &= \mathbb{E}[\mathbb{E}[dS'_\eta] (\mathbb{E}[S_\eta S'_\eta])^{-1} S_\eta S'_\eta (\mathbb{E}[S_\eta S'_\eta])^{-1} \mathbb{E}[S_\eta d']] \\ &= \mathbb{E}[d_\eta d'_\eta] \end{aligned}$$

Note that d_η is mean-zero since

$$\begin{aligned} \mathbb{E}[d_\eta] &= \mathbb{E}[\mathbb{E}[dS'_\eta] (\mathbb{E}[S_\eta S'_\eta])^{-1} S_\eta] \\ &= \mathbb{E}[dS'_\eta] (\mathbb{E}[S_\eta S'_\eta])^{-1} \mathbb{E}[S_\eta] \\ &= 0 \end{aligned}$$

This is useful because the Cramer-Rao bound of α under a parametric submodel equals the variance of d_η . Note further from the definition of d_η that it is the linear projection of pathwise-derivative d on score S_η . Therefore, the largest value of this projection can be obtained by considering the projection space as the scores corresponding to all parametric submodels. To formalize this, we next define the tangent set:

Definition 9 (Tangent Set). *Define the tangent set \mathcal{T} to be the mean square closure of all q -dimensional linear combinations of scores S_η for smooth parametric submodels:*

$$T = \{s \in \mathbb{R} : \mathbb{E}[\|s\|^2] \leq \infty, \quad \exists A_j S_{\eta_j} \quad \text{with} \quad \lim_{j \rightarrow \infty} \mathbb{E}[\|s - A_j S_{\eta_j}\|^2] = 0\}$$

The projection of d on the tangent set should have a larger variance than any particular submodel, suggesting that the projection should give the semi-parametric efficiency bound. The mathematical meaning of this projection on the tangent set is a least-squares projection in a Hilbert space of random vectors. This projection is defined as:

$$\delta \in \mathcal{T}, \quad \mathbb{E}[(d - \delta)s] = 0 \quad \text{for all} \quad s \in \mathcal{T}$$

If \mathcal{T} is linear, then δ exists and unique. It is called the efficient score because it equals the efficient influence function in asymptotically linear estimators.

Theorem 10 ([16], Theorem 3.1). *Suppose that the parameter is differentiable, \mathcal{T} is linear, and $\mathbb{E}[\delta\delta']$ is nonsingular, for the projection δ of d on \mathcal{T} . Then semi-parametric efficiency bound equals $\mathbb{E}[\delta\delta']$.*

D Proof of Theorem 2

Proof. We follow the steps outlined in Section (C.2) to calculate the semi-parametric efficiency bound of the parameter of interest:

$$\alpha := \mathbb{E}[\langle \theta_0(z), \phi(\pi(z), z) \rangle] \quad (48)$$

Let $f(y, a \mid z)$ and $f(z)$ denote the conditional distribution of (y, a) given z and the marginal distribution of z , respectively. The density of data (y, a, z) is then equal to:

$$f(y, a, z) = f(y, a \mid z)f(z)$$

We consider a regular parametric submodel, parameterized by η , to calculate the pathwise derivative of $\alpha(\eta)$:

$$f(y, a, z; \eta) = f(y, a \mid z; \eta)f(z; \eta)$$

The corresponding scores for the parametric submodel is given by:

$$s_\eta(y, a, z; \eta) = s_\eta(y, a \mid z; \eta) + s_\eta(z; \eta)$$

where $s_\eta(y, a, z; \eta) = 2 \frac{f_\eta(y, a, z; \eta)}{f(y, a, z; \eta)}$, and other scores are defined similarly.

Under the parametric submodel α can be written as:

$$\alpha(\eta) = \int \langle \theta(z; \eta), \phi(\pi(z), z) \rangle f(z; \eta) dz \quad (49)$$

The first step in semi-parametric efficiency bound derivation is to show that $\alpha(\eta)$ is pathwise differentiable, i.e. there exists $d(y, a, z; \eta_0)$ such that

$$\frac{\partial \alpha(\eta)}{\partial \eta} = \mathbb{E}[d(y, a, z; \eta) S_\eta(y, a, z; \eta)]$$

Let η_0 denote the true parameter value in the sense that $f(y, a, z; \eta_0)$ corresponds to the density of the data. To show pathwise differentiability, we differentiate Equation (49) under the integral sign and evaluate at $\eta = \eta_0$:

$$\frac{\partial \alpha(\eta_0)}{\partial \eta} = \int \left\langle \frac{\partial \theta(z; \eta_0)}{\partial \eta}, \phi(\pi(z), z) \right\rangle f(z; \eta_0) dz + \int \langle \theta_0(z; \eta_0), \phi(\pi(z), z) \rangle \frac{\partial f(z; \eta_0)}{\partial \eta} dz \quad (50)$$

$$= \mathbb{E} \left[\left\langle \frac{\partial \theta(z; \eta_0)}{\partial \eta}, \phi(\pi(z), z) \right\rangle \right] + \mathbb{E}[\langle \theta(z; \eta_0), \phi(\pi(z), z) \rangle s_\eta(z; \eta_0)] \quad (51)$$

To calculate $\partial\theta(z; \eta_0)/\partial\eta$ inside the expectations we use the first order conditions of the linear projection:

$$\begin{aligned}\mathbb{E}[(y - \langle\theta_0(z), \phi(a, z)\rangle)\phi(a, z) | z] &= 0 \\ \int (y - \langle\theta(z; \eta_0), \phi(a, z)\rangle)\phi_i(a, z)f(y, a | z; \eta_0)dyda &= 0\end{aligned}$$

Taking the derivative under the integral sign and evaluating at η_0 for all i :

$$\mathbb{E}\left[\left\langle\frac{\partial\theta(z; \eta_0)}{\partial\eta}, \phi(a, z)\phi(a, z)^T\right\rangle | z\right] + \mathbb{E}[(y - \langle\theta(z; \eta_0), \phi(a, z)\rangle)\phi(a, z)s_\eta(y, a | z, \eta_0) | z] = 0$$

Solving for $\partial\theta(z; \eta_0)/\partial\eta$

$$\partial\theta(z; \eta_0)/\partial\eta = \mathbb{E}[\Sigma(z)^{-1}\phi(a, z)(y - \langle\theta(z; \eta_0), \phi(a, z)\rangle)s_\eta(y, a | z; \eta_0) | z]$$

Substituting this into Equation (51):

$$\frac{\partial\alpha(\eta_0)}{\partial\eta} = \mathbb{E}[\langle\Sigma_0(z)^{-1}\phi(a, z)(y - \langle\theta_0(z), \phi(a, z)\rangle), \phi(\pi(z), z)\rangle s_\eta(y, a | z; \eta_0)] + \quad (52)$$

$$\begin{aligned}& \mathbb{E}[\langle\theta_0(z), \phi(\pi(z), z)\rangle s_\eta(z; \eta_0)] \\ &= \mathbb{E}[(\langle\theta_0(z) + \Sigma_0(z)^{-1}\phi(a, z)(y - \langle\theta_0(z), \phi(a, z)\rangle), \phi(\pi(z), z)\rangle - \alpha(\eta_0)) (s_\eta(y, a | z, \eta_0) + s_\eta(z; \eta_0))] \\ &= \mathbb{E}[d(y, a, z; \eta_0) (s_\eta(y, a | z; \eta_0) + s_\eta(z; \eta_0))] \\ &= \mathbb{E}[d(y, a, z; \eta_0) (s_\eta(y, a | z; \eta_0))] \quad (53)\end{aligned}$$

The second line follows because:

$$\mathbb{E}[\langle\theta_0(z), \phi(\pi(z), z)\rangle s_\eta(y, a | z, \eta_0)] = \mathbb{E}[\langle\theta_0(z), \phi(\pi(z), z)\rangle]\mathbb{E}[s_\eta(y, a | z, \eta_0) | z] = 0$$

$$\mathbb{E}[\alpha(\eta_0)s_\eta(z; \eta_0)] = \alpha(\eta_0)\mathbb{E}[s_\eta(z; \eta_0)] = 0$$

and

$$\mathbb{E}[\langle\Sigma_0(z)^{-1}\phi(a, z)(y - \langle\theta_0(z), \phi(a, z)\rangle), \phi(\pi(z), z)\rangle s_\eta(z; \eta_0)] = 0$$

Subtracting $\alpha(\eta_0)$ in the second line makes the pathwise derivative mean zero, which will prove useful later when projecting $d(y, a, z; \eta_0)$ on the tangent set.

Since Equation (52) satisfies the condition given in the definition of pathwise differentiability, the pathwise derivative of $\alpha(\eta)$ is:

$$d(y, a, z; \eta_0) = (\langle\theta_0(z) + \Sigma_0(z)^{-1}\phi(a, z)(y - \langle\theta_0(z), \phi(a, z)\rangle), \phi(\pi(z), z)\rangle - \alpha)$$

The semi-parametric efficiency bound for α is the variance of the projection of $d(y, a, z; \eta_0)$ onto the tangent space defined as the closed linear span of the scores:

$$\mathcal{T} = \{s(y, a | z) + s(z)\}$$

Note that the joint distribution is unrestricted so the only restrictions on the score functions are $E[s(y, x | z) | z] = 0$ and $E[s(z)] = 0$ and they are smooth.

Next, we show that the pathwise derivative is already in the tangent set $d(y, a, z; \eta_0) \in \mathcal{T}$. To see this we can write $d(y, a, z; \eta_0)$ as the sum of two functions:

$$d(y, a, z; \eta_0) = (\Sigma_0(z)^{-1}\phi(a, z)(y - \langle\theta_0(z), \phi(a, z)\rangle), \phi(\pi(z), z)\rangle) + (\langle\theta_0(z), \phi(\pi(z), z)\rangle - \alpha)$$

The first component is mean independent of z :

$$\mathbb{E}[\langle\Sigma_0(z)^{-1}\phi(a, z)(y - \langle\theta_0(z), \phi(a, z)\rangle), \phi(\pi(z), z)\rangle | z] = 0$$

The second component is function of only z and has zero mean:

$$\mathbb{E}[\langle\theta_0(z), \phi(\pi(z), z)\rangle - \alpha] = 0$$

Therefore, the pathwise derivative equals the sum of two functions that satisfy the restrictions on score functions in the tangent set, namely, $E[s(y, x | z) | z] = 0$ and $E[s(z)] = 0$. From this, we

conclude that $d(y, a, z; \eta_0)$ is in the tangent set; so the projection of $d(y, a, z; \eta_0)$ onto \mathcal{T} is equal to itself.

Therefore, the efficiency bound for α is:

$$\begin{aligned} V_{eff}(\alpha) &= Var(d(y, a, z; \eta_0)) \\ &= Var(v_{DR}(y, a, z; \pi)) \end{aligned}$$

Therefore, the doubly robust estimator, $v_{DR}(y, a, z; \pi)$, achieves the semi-parametric efficiency bound. This result extends to the difference of value functions by linearity of pathwise derivative.

To investigate the semi-parametric efficiency bound under the correct specification we use a result from [4] who shows that under the correct specification the efficiency bound is:

$$\begin{aligned} V_{eff}^c(\alpha) &= Var(\langle \theta_0(z), \phi(\pi(z), z) \rangle) \\ &\quad + \mathbb{E}[\phi(\pi(z), z) \mathbb{E}[\phi(a, z) \mathbb{E}[\epsilon^2 | a, z]^{-1} \phi(a, z)' | z]^{-1} \phi(\pi(z), z)^T] \end{aligned}$$

where $\epsilon = (y - \langle \theta_0(z), \phi(a, z) \rangle)$ is defined as residuals.

Under the homoskedasticity assumption, $\mathbb{E}[\epsilon^2 | a, z] = \sigma^2$, this efficiency bound becomes:

$$\begin{aligned} V_{eff}^c(\alpha) &= Var(\langle \theta_0(z), \phi(\pi(z), z) \rangle) + \\ &\quad \sigma^2 \mathbb{E}[\phi(\pi(z), z) \mathbb{E}[\phi(a, z) \phi(a, z)' | z]^{-1} \phi(\pi(z), z)^T] \\ &= Var(\langle \theta_0(z), \phi(\pi(z), z) \rangle) + \sigma^2 \mathbb{E}[\phi(\pi(z), z) \Sigma_0(z)^{-1} \phi(\pi(z), z)^T] \end{aligned}$$

which is equal to the variance of the doubly robust estimator:

$$\begin{aligned} V_{eff}(\alpha) &= Var(\langle \theta_0(z), \phi(\pi(z), z) \rangle) + \\ &\quad \mathbb{E}[\phi(\pi(z), z) \mathbb{E}[\Sigma_0(z)^{-1} \phi(a, z) \epsilon^2 \phi(a, z)' \Sigma(z)^{-1} | z] \phi(\pi(z), z)^T] \\ &= Var(\langle \theta_0(z), \phi(\pi(z), z) \rangle) + \\ &\quad \sigma^2 \mathbb{E}[\phi(\pi(z), z) \Sigma_0(z)^{-1} \mathbb{E}[\phi(a, z) \phi(a, z)' | z] \Sigma(z)^{-1} \phi(\pi(z), z)^T] \\ &= Var(\langle \theta_0(z), \phi(\pi(z), z) \rangle) + \\ &\quad \sigma^2 \mathbb{E}[\phi(\pi(z), z) \Sigma_0(z)^{-1} \Sigma_0(z)^{-1} \Sigma_0(z)^{-1} \phi(\pi(z), z)^T] \\ &= Var(\langle \theta_0(z), \phi(\pi(z), z) \rangle) + \sigma^2 \mathbb{E}[\phi(\pi(z), z) \Sigma_0(z)^{-1} \phi(\pi(z), z)^T] \\ &= Var(v_{DR}(y, a, z; \pi)) \end{aligned}$$

□

E Double Robustness Property of Policy Estimator

Theorem 11 (Double Robustness). $V_{DR}(\pi)$ is an unbiased estimate of $V_0(\pi(z), z)$ if for all z , either $\mathbb{E}_{S_1 \sim D^{n/2}}[\hat{\theta}(z)] = \theta_0(z)$ or $\mathbb{E}_{S_1 \sim D^{n/2}}[\hat{\Sigma}(z)^{-1}] = \Sigma_0(z)^{-1}$, where expectation is taken over the randomness of the nuisance estimation sample S_1 .

Proof. Let $\bar{\theta}(z) = \mathbb{E}_{S_1 \sim D^{n/2}}[\hat{\theta}(z)]$ and $\bar{\Sigma}^{-1}(z) = \mathbb{E}_{S_1 \sim D^{n/2}}[\hat{\Sigma}(z)^{-1}]$, be the expected value of the estimates at any input z , where the expectation is with respect to the randomness on the half-split of $n/2$ samples that were used for training the estimates. Due to sample-splitting and cross-fitting, the expected value of the doubly robust policy estimate can be written as:

$$\mathbb{E}[V_{DR}(\pi)] = \mathbb{E}[\langle \bar{\theta}(z) + \bar{\Sigma}(z)^{-1} \phi(a, z) (y - \langle \bar{\theta}(z), \phi(a, z) \rangle), \phi(\pi(z), z) \rangle] \quad (54)$$

where the random variables (y, a, z) are a fresh independent draw of the data generating process that generated the observational data.

Observe that y is an unbiased estimate of $V(a, z)$ conditional on z . Moreover, since $\theta_0(z)$ is the minimizer of the conditional squared loss, taking the first order condition implies:

$$\begin{aligned} \mathbb{E}[(V_0(a, z) - \langle \theta_0(z), \phi(a, z) \rangle) \phi(a, z) | z] &= 0 \iff \\ \mathbb{E}[y \phi(a, z) | z] &= \mathbb{E}[\langle \theta_0(z), \phi(a, z) \rangle \phi(a, z) | z] \end{aligned}$$

Thus we can re-write the expected value of the doubly robust policy estimate as:

$$\begin{aligned}
\mathbb{E}[V_{DR}(\pi)] &= \mathbb{E}[\langle \bar{\theta}(z) + \bar{\Sigma}(z)^{-1} \phi(a, z) (Y - \langle \bar{\theta}(z), \phi(a, z) \rangle), \phi(\pi(z), z) \rangle] \\
&= \mathbb{E}[\langle \bar{\theta}(z) + \bar{\Sigma}(z)^{-1} \phi(a, z) \langle \theta_0(z) - \bar{\theta}(z), \phi(a, z) \rangle, \phi(\pi(z), z) \rangle] \\
&= \mathbb{E}[\langle \bar{\theta}(z) + \bar{\Sigma}(z)^{-1} \phi(a, z) \phi(a, z)^T (\theta_0(z) - \bar{\theta}(z)), \phi(\pi(z), z) \rangle] \\
&= \mathbb{E}[\langle \bar{\theta}(z) + \bar{\Sigma}(z)^{-1} \mathbb{E}[\phi(a, z) \phi(a, z)^T | z] (\theta_0(z) - \bar{\theta}(z)), \phi(\pi(z), z) \rangle] \\
&= \mathbb{E}[\langle \bar{\theta}(z) + \bar{\Sigma}(z)^{-1} \Sigma_0(z) (\theta_0(z) - \bar{\theta}(z)), \phi(\pi(z), z) \rangle] \\
&= \mathbb{E}[\langle \bar{\theta}(z) + \bar{\Sigma}(z)^{-1} \Sigma_0(z) (\theta_0(z) - \bar{\theta}(z)), \phi(\pi(z), z) \rangle]
\end{aligned}$$

Hence we have:

$$\mathbb{E}[V_{DR}(\pi)] - V_0(\pi) = \mathbb{E}[\langle (\bar{\Sigma}(z)^{-1} \Sigma_0(z) - \mathbb{I}) (\theta_0(z) - \bar{\theta}(z)), \phi(\pi(z), z) \rangle]$$

The right hand side is zero if either $\bar{\theta}(z) = \theta_0(z)$ or if $\bar{\Sigma}(z)^{-1} = \Sigma_0(z)$. \square

F Lipschitz Variogram Settings and Binary Treatment

For simplicity of notation, we let $v(x; \pi) = v_{DR}^0(x; \pi)$ and $\pi_* = \pi_*^0$ throughout this section, as the results are not specific to the doubly robust value function. Suppose that the value function of the policy learning problem has the following self-bounded Lipschitz property:

$$\text{Var}(v(x; \pi)) - C \text{Var}(v(x; \pi_*)) \leq L |\mathbb{E}[v(x; \pi)] - \mathbb{E}[v(x; \pi_*)]| = L(V(\pi_*) - V(\pi))$$

for some constants C, L , i.e. if a policy has value close to the optimal policy, then it does not have much larger variance. Then we have that:

$$\begin{aligned}
\sup_{\pi, \pi' \in \Pi_*(\epsilon_n)} \text{Var}(v(x; \pi) - v(x; \pi')) &\leq \sup_{\pi \in \Pi_*(\epsilon_n)} 4 \text{Var}(v(x; \pi)) \\
&\leq 4C \text{Var}(v(x; \pi_*)) + 4L \sup_{\pi \in \Pi_*(\epsilon_n)} (V(\pi) - V(\pi_*)) \\
&\leq 4C \underbrace{\text{Var}(v(x; \pi_*))}_{V_*} + 4L \epsilon_n
\end{aligned}$$

Thus we get regret rates of the form:

$$\begin{aligned}
V(\pi_*) - V(\pi_2) &= O\left(\kappa(2\sqrt{CV_*}, \mathcal{F}_\Pi) + \sqrt{\frac{V_* \log(1/\delta)}{n}} + \epsilon_n \frac{1}{\sqrt{n}}\right) \\
&= O\left(\kappa(2\sqrt{CV_*}, \mathcal{F}_\Pi) + \sqrt{\frac{V_* \log(1/\delta)}{n}}\right)
\end{aligned}$$

since $\epsilon_n = o(1)$.

Example 3 (Binary Treatment). *In the case of binary treatment considered in [1], the loss took the form:*

$$v(x; \pi) = \Gamma(z) \cdot (2\pi(z) - 1) \quad (55)$$

with $\pi : Z \rightarrow \{0, 1\}$. In this case observe that the self-bounded property is satisfied since:

$$\begin{aligned}
\text{Var}(v(x; \pi)) &= \mathbb{E}[v(x; \pi)^2] - \mathbb{E}[v(x; \pi)]^2 \\
&= \mathbb{E}[\Gamma(z)^2 (2\pi(z) - 1)^2] - V(\pi)^2 \\
&= \mathbb{E}[\Gamma(z)^2] - V(\pi)^2
\end{aligned}$$

Where the latter property holds since $(2\pi(z) - 1)^2 = 1$ irrespective of $\pi(z)$. Thus the first part in the variance is independent of the policy, which is the crucial special property of the binary treatment case. This leads to the fact that:

$$\text{Var}(v(x; \pi)) - \text{Var}(v(x; \pi_*)) = V(\pi_*)^2 - V(\pi)^2 \leq 2 |V(\pi) - V(\pi_*)| \quad (56)$$

Hence, the self-boundedness property holds with $C = 1$ and $L = 2$. Thus for the binary treatment setting we can achieve a regret rate whose leading term only depends on the semi-parametric efficient variance of the optimal policy.

As a concrete example, consider the case when the class \mathcal{F}_Π is a VC-subgraph class of VC dimension d , and let $S_n = \mathbb{E}_n[\sup_\pi v(x; \pi)^2] = \mathbb{E}_n[\Gamma(z)^2]$. Then Theorem 2.6.7 of [22] shows that: $\mathcal{H}_2(\epsilon, \mathcal{F}_\Pi, n) = O(d(1 + \log(S_n/\epsilon)))$. This implies that

$$\kappa(r, \mathcal{F}_\Pi) = O\left(\int_0^r \sqrt{d(1 + \log(S_n/\epsilon))} d\epsilon\right) = O\left(r\sqrt{d}\sqrt{1 + \log(S/r)}\right).$$

Moreover, by Markov's inequality w.p. $1 - \delta$, $S_n \leq \mathbb{E}[S_n]/\delta = \mathbb{E}[\sup_\pi v(x; \pi)^2]/\delta = \mathbb{E}[\Gamma(z)^2]/\delta := S/\delta$. Hence, we can conclude that w.p. $1 - \delta$:

$$V(\pi_*) - V(\pi_2) = O\left(r\sqrt{1 + \log(S/r)}\sqrt{\frac{d}{n}} + r\sqrt{\frac{\log(1/\delta)}{n}} + \frac{d(1 + \log(S/r))\log(1/\delta)}{n}\right).$$

Combining all the above we get a bound of the form (excluding lower order terms):

$$V(\pi_*) - V(\pi_2) = O\left(\sqrt{V_*(1 + \log(S/V_*))}\sqrt{\frac{d}{n}} + \sqrt{\frac{V_* \log(1/\delta)}{n}}\right).$$

which recovers the result of [1] up to constants.

G Doubly Robust Estimators in Pricing Experiment

G.1 Linear Model

We want to estimate some regression models of $a(z)$ and $b(z)$ in the demand model. For instance, if these fall in some high-dimensional linear function class, we can estimate a regression between demand and the linear function class. Moreover, we need to estimate the covariance matrix, which in this case takes the simple form:

$$\Sigma_0(z) = \begin{bmatrix} 1 & \mathbb{E}[p | z] \\ \mathbb{E}[p | z] & \mathbb{E}[p^2 | z] \end{bmatrix} \quad (57)$$

whose inverse takes the form:

$$\Sigma_0(z)^{-1} = \frac{1}{\text{Var}(p | z)} \begin{bmatrix} \mathbb{E}[p^2 | z] & -\mathbb{E}[p | z] \\ -\mathbb{E}[p | z] & 1 \end{bmatrix} \quad (58)$$

If for instance the observational policy was homoskedastic (i.e. the exploration component was independent of the context z), then $\text{Var}(p | z)$ is a constant σ^2 independent of z . Moreover, we can write:

$$\mathbb{E}[p^2 | z] = \sigma^2 + \mathbb{E}[p | z]^2 \quad (59)$$

Thus we only need to estimate the mean treatment policy $g(z) = \mathbb{E}[p | z]$ and the variance σ^2 . Then the doubly robust estimate of $a(z)$ takes the form:

$$\begin{aligned} a_{DR}(z) &= \hat{a}(z) + \left(1 + \hat{g}(z)\frac{\hat{g}(z) - p}{\hat{\sigma}^2}\right) (d - \hat{a}(z) - \hat{b}(z)p) \\ b_{DR}(z) &= \hat{b}(z) + \frac{p - \hat{g}(z)}{\hat{\sigma}^2} (d - \hat{a}(z) - \hat{b}(z)p) \end{aligned}$$

G.2 Quadratic Model

In the case where we observe the revenue our model becomes quadratic in prices

$$r = a(z)x - b(z)x^2 + \epsilon$$

The covariance matrix takes the form:

$$\Sigma_0(z) = \begin{bmatrix} \mathbb{E}[p^2 | z] & \mathbb{E}[p^3 | z] \\ \mathbb{E}[p^3 | z] & \mathbb{E}[p^4 | z] \end{bmatrix}$$

whose inverse is:

$$\Sigma_0(z)^{-1} = \frac{1}{\mathbb{E}[p^4 | z]\mathbb{E}[p^2 | z] - \mathbb{E}[p^3 | z]^2} \begin{bmatrix} \mathbb{E}[p^4 | z] & -\mathbb{E}[p^3 | z] \\ -\mathbb{E}[p^3 | z] & \mathbb{E}[p^2 | z] \end{bmatrix}$$

Let $\mu_k(z)$ denote $E[p^k | z]$. If the observational policy was homoskedastic and none of the central moments of price depends on z , using the recursive structure, the nuisance functions in the covariance matrix can be written as

$$\begin{aligned}\mu_2(z) &= \mu_2^c + \mu_1(z)^2 \\ \mu_3(z) &= \mu_3^c + 3\mu_2(z)\mu_1(z) - 2\mu_1(z)^3 \\ \mu_4(z) &= \mu_4^c + 4\mu_3(z)\mu_1(z) - 6\mu_1(z)\mu_2(z) + 3\mu_1(z)^4\end{aligned}$$

where μ_k^c denotes the k -th central moment of p . Therefore, we only need to estimate the mean treatment policy $\mu_1(z)$ and the central moments μ_2^c , μ_3^c and μ_4^c . Then, the doubly robust estimate of $a(z)$ and $b(z)$ take the form:

$$\begin{aligned}a_{DR}(z) &= \hat{a}(z) + \left(\frac{\mu_4(z)p - \mu_3(z)p^2}{\mu_4(z)\mu_2(z) - \mu_3(z)^2} \right) (d - \hat{a}(z)p - \hat{b}(z)p^2) \\ b_{DR}(z) &= \hat{b}(z) + \left(\frac{\mu_2(z)p^2 - \mu_3(z)p}{\mu_4(z)\mu_2(z) - \mu_3(z)^2} \right) (d - \hat{a}(z)p - \hat{b}(z)p^2)\end{aligned}$$

H Additional Experiment Results

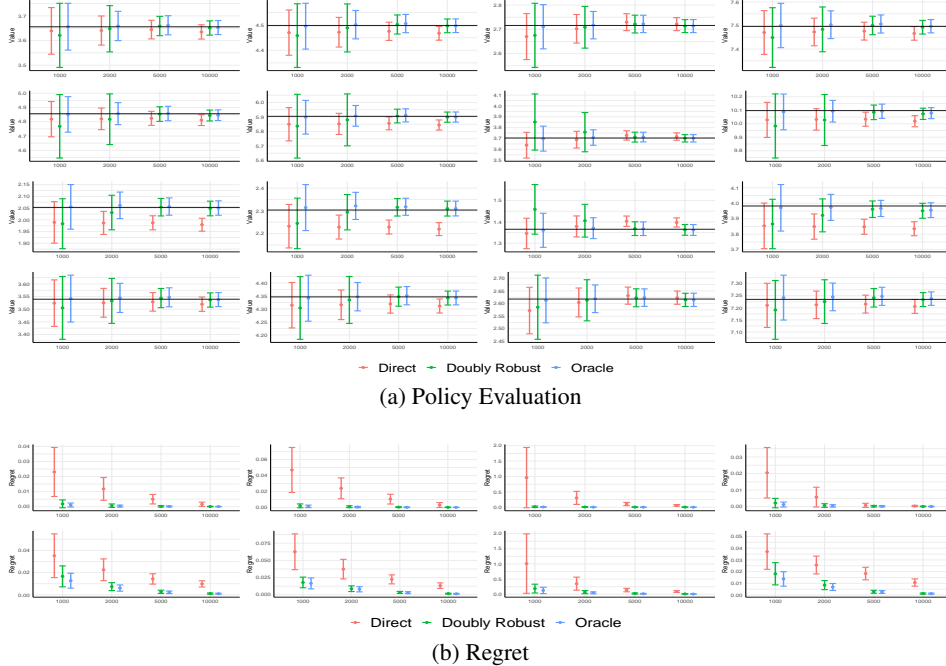
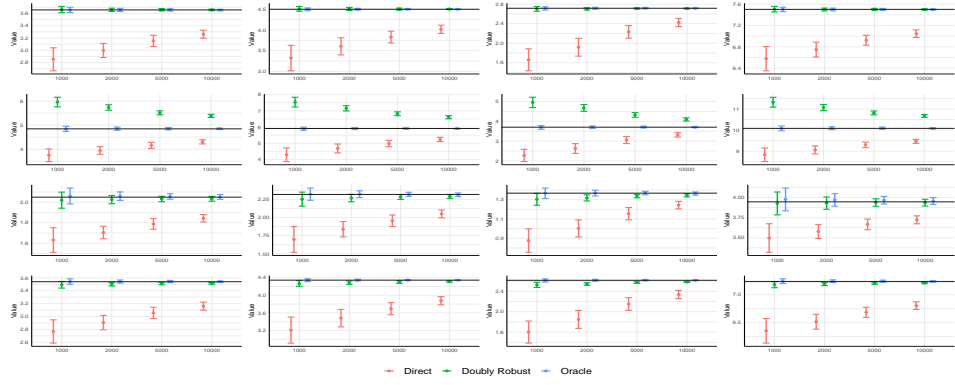


Figure 4: Linear, High Dimensional Regime: (a) Black line shows the true value of the policy, and each line shows the mean and standard deviation of the policy over 100 simulations. (b) each line shows the mean and standard deviation of the value of the corresponding policy over 100 simulations. We omit the results for the inverse propensity score method since they are too large to report together with the other estimates in the high dimensional regime.



(a) Policy Evaluation

Figure 5: Quadratic, High Dimensional Regime: (a) Black line shows the true value of the policy, and each line shows the mean and standard deviation of the policy over 100 simulations. (b) each line shows the mean and standard deviation of the value of the corresponding policy over 100 simulations. We omit the results for the inverse propensity score method since they are too large to report together with the other estimates in the high dimensional regime.