
Flexible and Efficient Contextual Bandits with Heterogeneous Treatment Effect Oracles

Aldo Gael Carranza*
Stanford University

Sanath Kumar Krishnamurthy*
Stanford University

Susan Athey
Stanford University

Abstract

Contextual bandit algorithms often estimate reward models to inform decision-making. However, true rewards can contain action-independent redundancies that are not relevant for decision-making. We show it is more data-efficient to estimate any function that explains the reward differences between actions, that is, the *treatment effects*. Motivated by this observation, building on recent work on oracle-based bandit algorithms, we provide the first reduction of contextual bandits to general-purpose heterogeneous treatment effect estimation, and we design a simple and computationally efficient algorithm based on this reduction. Our theoretical and experimental results demonstrate that heterogeneous treatment effect estimation in contextual bandits offers practical advantages over reward estimation, including more efficient model estimation and greater flexibility to model misspecification.

1 INTRODUCTION

We study cumulative regret minimization in the stochastic contextual bandit setting (Li et al., 2010). A common approach in contextual bandits is to use a specified function class to estimate a model of the expected reward conditional on any given context-action pair and utilize the estimated model to recommend actions (Chu et al., 2011; Agrawal and Goyal, 2013; Foster and Rakhlin, 2020; Simchi-Levi and Xu, 2021). Such a model estimate is always biased when the given model class does not capture the true conditional expected reward (e.g., a linear model class is used in reward estimation when the true reward model is nonlinear). This bias results in an unavoidable

additional linear regret term, which is asymptotically dominant, in the corresponding contextual bandit algorithm (Krishnamurthy et al., 2021; Foster et al., 2020).

In this paper, we obtain tighter cumulative regret bounds by reducing the linear term due to model bias by shifting the target of model estimation for decision-making. We observe that the conditional expected reward does not need to be entirely estimated to make decisions. It suffices to estimate any function that explains the reward differences between actions, that is, any model that captures *treatment effects*. We show the bias of treatment effect estimation is no greater (and often significantly smaller) than the bias of a reward model estimation under the same model class. An example of lower bias in treatment effect estimation for a fixed model class is the case where the true reward is the sum of a non-linear function of only contexts and a linear function of contexts and actions. In this case, while the conditional reward functions may require complex model classes to be well approximated, a linear function class can effectively capture the true treatment effects.

Motivated by these observations, we design an improved contextual bandit algorithm that relies on an oracle for heterogeneous treatment effect (HTE) estimation known as the *R*-learner (Nie and Wager, 2021). The *R*-learner possesses certain properties that make it more useful as a heterogeneous treatment effect estimator for model estimation in contextual bandits compared to other comparable methods, including *S*-learner, *T*-learner, or *X*-learner (Künzel et al., 2019). The principal reason is that, unlike other methods which attempt to form uniformly accurate estimates of treatment effects over the context distribution, the *R*-learner estimates treatment effects more accurately in regions of exploration in the context space where actions are sampled non-deterministically. Since contextual bandits explore more in regions of context space where there is more uncertainty regarding the best action, the *R*-learner’s more accurate treatment effect estimates in these regions allow our algorithms to achieve optimal regret guarantees.

In this paper, we provide the first general-purpose reduction of contextual bandits to heterogeneous treatment effect es-

timization via R -loss minimization. Using this construction, we design a computationally efficient algorithm for contextual bandits with an R -loss oracle. We show that our approach is more robust to model misspecification than reward estimation methods based on squared error regression oracles. Experimentally, extending our approach to consider model selection, we validate that our algorithm outperforms these regression-based methods in settings where the treatment effect lies in a simpler class of functions than that of the reward.

1.1 Related Work

Contextual Bandits Contextual bandit algorithms can generally be classified into two categories of realizability-based or agnostic approaches. Realizability-based methods (Filippi et al., 2010; Abbasi-Yadkori et al., 2011; Chu et al., 2011; Agrawal and Goyal, 2013; Foster et al., 2018; Foster and Rakhlin, 2020; Simchi-Levi and Xu, 2021) are computationally simple yet require stricter assumptions on the underlying reward model, and agnostic methods (Auer et al., 2002; Langford and Zhang, 2007; Dudik et al., 2011; Beygelzimer et al., 2011; Agarwal et al., 2014) have weaker assumptions but are generally computationally intractable (Bietti et al., 2021). Our approach attempts to achieve a middle ground that captures the simplicity of realizability-based methods and robustness of agnostic methods. Our work extends a recent line of research on contextual bandit algorithms with least squares regression oracles (Foster and Rakhlin, 2020; Simchi-Levi and Xu, 2021) that utilize the inverse gap weighting action-sampling strategy of Abe and Long (1999). In contrast to this work, we develop a more flexible model estimation strategy built on a heterogeneous treatment effect estimation oracle primitive. Foster et al. (2020) and Krishnamurthy et al. (2021) extend the least squares regression oracle-based approach to the setting where reward model realizability may not hold and the misspecification error is unknown. Our approach relies on the misspecification adaptation strategies presented in these works.

Similar to our approach, Greenewald et al. (2017); Krishnamurthy et al. (2018); Peng et al. (2019); Choi et al. (2022) deal with estimating the essential components of a reward function for decision making. Although Krishnamurthy et al. (2018) use a similar local centering transformation as the R -learner to estimate treatment effects, their UCB-style approach is only applicable for linear treatment effect function classes. Similarly, Greenewald et al. (2017); Peng et al. (2019); Choi et al. (2022) consider UCB-based or Thompson sampling-based approaches to estimate a partially linear model. In contrast, our approach is flexible for any specified class of treatment effect models, not just linear models.

Various works have taken a similar yet distinct approach

on expanding reward estimation in contextual bandits beyond the standard setting. The problem of non-stationary confounding in contextual bandits is considered in Qin and Russo (2022); Choi et al. (2022), it is addressed using a Thompson sampling-based approach. Our algorithm provides an oracle-based approach that is capable of handling similar issues of non-stationary confounding without the restrictive assumptions of Thompson-sampling (see Section 6). Hu et al. (2020); Gur et al. (2022) consider contextual bandits with non-parametric reward functions with unknown smoothness which allows for more flexible model estimation in a general class of problems. Our work is supplemented by this work on non-parametric contextual bandits since it is also capable of handling non-parametric model classes for estimating the underlying treatment effect model. Contextual dueling bandits (Dudík et al., 2015; Saha and Krishnamurthy, 2022) consider the problem where the learner only observes the reward differences between pairs of actions. In contrast, although our approach captures residual reward differences from the mean, we are still in the standard setting where reward outcomes from individually selected arms are observed.

Heterogeneous Treatment Effect Estimation There are plenty of heterogeneous treatment effect estimation methods in the literature including the S -learner, T -learner, X -learner, and DR-learner (Künzel et al., 2019; Kennedy, 2020). As discussed in the introduction, we found the R -learner satisfies certain properties that make it a more natural fit for oracle-based contextual bandits. The R -learner was proposed in its full generality by Nie and Wager (2021) as a flexible and general-purpose method for heterogeneous treatment effect estimation amenable to machine learning methods since it is formulated as a risk minimization procedure over a specified class of treatment effect models. It is based on Robinson’s local centering transformation (Robinson, 1988) for estimating parametric components of partially linear models, which recent methods have built on for more general semi-parametric inference of treatment effects (Athey et al., 2019; Chernozhukov et al., 2017).

The exploration parameters of our algorithm rely on knowledge of excess risk bounds on the R -learner. Nie and Wager (2021) and Foster and Syrgkanis (2019) provide excess risk bounds on the R -learner which are not applicable to our setting. We look towards existing results on problem-dependent error bounds in Koltchinskii (2011); Xu and Zeevi (2020) that we show are applicable for establishing finite-sample excess risk bounds for the multi-treatment R -learner with known propensities.

2 PRELIMINARIES

Learning Setting Let \mathcal{A} be a set of K actions and \mathcal{X} be a space of contexts. A sequence of interactions between a learner and nature occurs over T rounds, where

T is possibly unknown. At each round t , nature randomly samples a context-reward pair (x_t, r_t) from a fixed but unknown probability distribution D over the joint space of contexts and reward vectors $\mathcal{X} \times [0, 1]^K$. The learner observes x_t , chooses an action $a_t \in \mathcal{A}$, and receives the reward $r_t(a_t) \in [0, 1]$ for the chosen action. Note that the rewards are assumed to be within $[0, 1]$ for simplicity, but our results can easily be extended more generally to bounded rewards.

Define the true conditional mean reward function to be $f^*(x, a) := \mathbb{E}[r(a)|x]$ for any $(x, a) \in \mathcal{X} \times \mathcal{A}$, where the expectation is taken with respect to D . The policy that maximizes the true conditional mean reward is $\pi^*(x) := \arg \max_a f^*(x, a)$. The goal of the learner is compete with the optimal policy π^* to minimize cumulative regret:

$$\text{Reg}(T) := \sum_{t=1}^T (r_t(\pi^*(x_t)) - r_t(a_t)). \quad (1)$$

Reward Decomposition Note that, for any function h^* of only contexts, there is a corresponding function g^* of contexts and actions such that

$$f^*(x, a) = g^*(x, a) + h^*(x) \quad (2)$$

for all $(x, a) \in \mathcal{X} \times \mathcal{A}$. Under any such decomposition of the conditional expected reward, we will refer to h^* as a *confounder* model and g^* as the corresponding *true treatment effect* model. A true treatment effect model is named as such because it sufficiently captures the conditional reward differences between actions, that is, given any $x \in \mathcal{X}$ we have that $g^*(x, a) - g^*(x, a') = f^*(x, a) - f^*(x, a')$ for any $a, a' \in \mathcal{A}$.

Note that such a decomposition of the conditional reward function is not unique and, in fact, can be trivial, i.e., $h^* \equiv 0$. However, if there exists a decomposition such that g^* can be captured by a ‘‘simpler’’ model class than that required to capture f^* , the learner could more efficiently estimate g^* instead for decision-making. This would be sufficient since the optimal policy can be equivalently defined as $\pi^*(x) = \arg \max_a g^*(x, a)$.

We assume that the learner has access to a class \mathcal{G} of measurable functions $g : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$ that models a true treatment effect. These functions are referred to as *treatment effect* models. We do not require treatment effect model realizability, i.e., there does not necessarily exist a model $g \in \mathcal{G}$ such that $g(x, a) = g^*(x, a)$ for all $(x, a) \in \mathcal{X} \times \mathcal{A}$. We discuss treatment effect model misspecification in Section 3.

Additional Notation The probability simplex over actions \mathcal{A} is denoted as $\Delta_{\mathcal{A}} := \{q \in [0, 1]^K \mid \sum_{a \in \mathcal{A}} q_a = 1\}$. For any probability kernel $p : \mathcal{X} \rightarrow \Delta_{\mathcal{A}}$ in the set of all probability kernels, denoted as \mathcal{P} , we let $p(\cdot|x)$ denote the associated probability distribution over \mathcal{A} for a

given context $x \in \mathcal{X}$. We denote $D(p)$ to be the observable data-generating distribution over $\mathcal{X} \times \mathcal{A} \times [0, 1]$ induced by $p \in \mathcal{P}$ on the distribution D such that sampling $(x, a, r(a)) \sim D(p)$ is defined as sampling $(x, r) \sim D$ then sampling $a \sim p(\cdot|x)$ and only revealing $(x, a, r(a))$. We let $D_{\mathcal{X}}$ denote the marginal distribution of D on the set of contexts \mathcal{X} . With a slight abuse of notation, for any function $g : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$ and context $x \in \mathcal{X}$, we let $g(x, \cdot)$ denote the vector $(g(x, a))_{a \in \mathcal{A}}$. Let $\langle \cdot, \cdot \rangle$ denote the inner product operator.

3 HETEROGENEOUS TREATMENT EFFECT ORACLE

We will use the R -learner, a method for heterogeneous treatment effect estimation based on R -loss minimization (Nie and Wager, 2021), as the model fitting subroutine in our contextual bandit algorithm. In this section, we describe the multi-treatment R -loss and the R -learner model fitting oracle we assume access to. We also discuss the benefits of the R -learner for contextual bandits.

3.1 R -loss Risk, Excess Risk, and Misspecification Error

Let $p : \mathcal{X} \rightarrow \Delta_{\mathcal{A}}$ be a fixed probability kernel and let $e_a \in \{0, 1\}^K$ be the one-hot encoding of action $a \in \mathcal{A}$. The multi-treatment R -loss *risk* of a model $g \in \mathcal{G}$ under kernel p is

$$\mathcal{R}_p(g) := \mathbb{E} \left[(r(a) - \hat{\mu}(x) - \langle e_a - p(x), g(x, \cdot) \rangle)^2 \right], \quad (3)$$

where the expectation is taken with respect to $D(p)$ and $\hat{\mu}$ is an estimate of the conditional mean of rewards realized under the action-sampling kernel p , defined as $\mu(x) := \mathbb{E}_{a \sim p(\cdot|x)}[f^*(x, a)]$ for all $x \in \mathcal{X}$.¹

The R -loss *excess risk* of a model $g \in \mathcal{G}$ under kernel p is

$$\mathcal{E}_p(g) := \mathcal{R}_p(g) - \min_{g'} \mathcal{R}_p(g'), \quad (4)$$

where the infimum is taken with respect to all functions. The R -loss excess risk has an important role in our analysis since it is fundamental to the evaluation of treatment effect model quality.

The *average misspecification error* of the model class \mathcal{G} under the R -loss is defined as the quantity \sqrt{B} ,² where

$$B := \max_{p \in \mathcal{P}} \min_{g \in \mathcal{G}} \mathcal{E}_p(g). \quad (5)$$

Capturing this quantity of model misspecification is useful for doing away with an assumption on model realizability.

¹We use an estimate of μ in our definition of R -loss risk since it is unknown in our problem setting. The action-sampling probabilities under kernel p , on the other hand, are known.

²Similar measures of misspecification are denoted by ϵ in other works on misspecified models (Ghosh et al., 2017; Foster and Rakhlin, 2020; Foster et al., 2020; Lattimore et al., 2020).

3.2 R -loss Properties

The R -loss is based on Robinson’s local centering transformation³ and was introduced as a minimization objective for flexible heterogeneous treatment effect estimation (Robinson, 1988; Nie and Wager, 2021). Indeed, the infima of the R -loss risk correspond exactly with the true treatment effect models (see Appendix A.1 for further discussion).

The nuisance parameters μ and p of the R -loss are typically estimated from data. In our setting, the exact treatment propensities $p(a|x)$ are known to the bandit learner because they are exactly the action-sampling probabilities of the bandit algorithm and these quantities suffice for model estimation. However, the conditional mean realized rewards $\mu(x)$ are not known, so μ must be estimated. Even so, the impact of this nuisance parameter estimation is almost negligible to our analysis as we show in the following proposition that the R -loss achieves a form of robustness to unknown nuisance parameter μ in the sense that the excess R -loss risk does not depend on any fixed estimate of μ when p is known. See Appendix A.1 for the proof.

Proposition 1. *Suppose $\hat{\mu}$ is any fixed function. For any $g^* \in \arg \min_{g'} \mathcal{R}_p(g')$,*

$$\mathcal{E}_p(g) = \mathbb{E} \left[\langle e_a - p(x), g(x, \cdot) - g^*(x, \cdot) \rangle^2 \right], \quad (6)$$

where the expectation is taken with respect to $D(p)$.

This result holds for any given model $g \in \mathcal{G}$, and so any uniform estimate $\hat{\mu}$ suffices (e.g., $\hat{\mu} \equiv 0$). However, a more accurate estimate of μ reduces the variance of the R -loss,⁴ thus improving the finite-sample excess risk rates of estimating the treatment effect model by a problem-dependent constant factor. Therefore, in our implementation, we choose to train $\hat{\mu}$ on available data using a *cross-fitting* strategy where the function is estimated on different data folds and the output $\hat{\mu}(x)$ for a given point x is the output under the function estimated on the fold corresponding to x . This ensures a better estimate of the nuisance parameter that is still fixed relative to the sampled data, as suggested in Chernozhukov et al. (2017); Nie and Wager (2021).

Note that Equation 6 vanishes when the action-sampling kernel is deterministic. Therefore, we observe that Proposition 1 reveals that the R -loss excess risk does not penalize models with less accurate treatment effect estimates in regions where actions are sampled deterministically. This

³The multivariate version of Robinson’s transformation is $r(a) - \mu(x) = \langle e_a - p(x), g^*(x, \cdot) \rangle + \varepsilon$ where ε is a noise term such that $\mathbb{E}[\varepsilon|x, a] = 0$.

⁴In empirical risk minimization, excess risk rates depend on both the variance of the loss being minimized and the complexity of the hypothesis class. Due to randomness in sampled data, loss variance is typically non-zero.

allows capturing a more precise form of uncertainty quantification useful for proving tighter regret bounds. In particular, this allows us to prove the following proposition showing that, for a given model class, the average misspecification error under the R -loss is bounded above by a corresponding notion of average misspecification error under the *squared error loss* typically used in reward model estimation (Simchi-Levi and Xu, 2021).

Proposition 2. *For a fixed model class \mathcal{G} , the average misspecification error under the R -loss is no greater than the average misspecification error under the squared error loss, i.e.,*

$$B \leq \max_{p \in \mathcal{P}} \min_{f \in \mathcal{G}} \mathbb{E} \left[(f(x, a) - f^*(x, a))^2 \right]. \quad (7)$$

where the expectation is taken with respect to $D(p)$.

This result will be important in highlighting the benefits of treatment effect estimation over reward estimation in ensuring tighter regret bounds under model misspecification. See Appendix A.2 for a proof and further discussion on the intuition of this result.

3.3 R -loss Oracle

We will assume access to an oracle that minimizes the R -loss risk given finite data samples such that the excess risk of the oracle output can be bounded by a known estimation error rate on the given model class \mathcal{G} and the misspecification error B . We present this formally in the following assumption.

Assumption 1. *Let \mathcal{D} be a data set of n independently and identically drawn samples from $D(p)$ under some fixed kernel p . Let \hat{g} be the output model estimate of a given R -loss oracle with \mathcal{D} as input. For any $\zeta \in (0, 1)$, the R -loss excess risk of \hat{g} is bounded as*

$$\mathcal{E}_p(\hat{g}) \leq \xi(n, \zeta) + C_0 B, \quad (8)$$

with probability $1 - \zeta$, where C_0 is a universal constant and $\xi(n, \zeta)$ is the known estimation rate for the model class \mathcal{G} that is non-increasing in $n \in \mathbb{N}$ for all $\zeta \in (0, 1)$.

An example of an estimator that satisfies Assumption 1 is the R -learner which performs empirical R -loss risk minimization over the specified class of models \mathcal{G} :

$$\hat{g} \in \arg \min_{g \in \mathcal{G}} \widehat{\mathcal{R}}_p(g; \mathcal{D}) + \Lambda(g), \quad (9)$$

where $\mathcal{D} = \{(x_t, a_t, r_t(a_t))\}_{t \in \mathcal{T}}$ is the given input dataset indexed by \mathcal{T} of samples drawn i.i.d. from $D(p)$, Λ is a regularizer on the complexity of the treatment effect model, and

$$\widehat{\mathcal{R}}_p(g; \mathcal{D}) = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \left(r_t(a_t) - \hat{\mu}_t(x_t) - \langle e_{a_t} - p(x_t), g(x_t, \cdot) \rangle \right)^2 \quad (10)$$

is the empirical R -loss risk which uses an estimate $\hat{\mu}_t$ for μ that is constructed to be independent of $(x_t, a_t, r_t(a_t))$ for each $t \in \mathcal{T}$, typically via cross-fitting as previously discussed.

In Appendix A.3, we discuss valid finite-sample R -loss excess risk rates for the R -learner under various important classes of models including classes with finite VC dimension and non-parametric classes of polynomial growth.⁵

Remark 1. Note that the oracle assumption as stated requires i.i.d. data from the data-generating process induced by a fixed kernel. In our bandit setting where model estimates inform kernel updates, this requirement would seem to imply that model estimation must be done only using data collected with the most recent kernel and any update on the kernel requires a subsequent refresh on the data buffer, reducing the effective sample size. However, this i.i.d. requirement is purely due to technical reasons and we can do away with it using martingale arguments that extend our analysis to adaptive kernels. Nevertheless, we keep the i.i.d. assumption in this paper to maintain focus on our contributions to more efficient model estimation via treatment effect estimation and to avoid a more complicated discussion of martingales.

3.4 R -loss Benefits for Contextual Bandits

We elaborate on the benefits of R -learner over other heterogeneous treatment effect estimators for the contextual bandit setting. For simplicity, we restrict our discussion to the two arm setting $\mathcal{A} = \{a_0, a_1\}$. This is only a simplification for the purposes of this discussion as our algorithm allows for more than two arms.

In this setting, estimating the true HTE function defined by $\tau^*(x) := f^*(x, a_1) - f^*(x, a_0)$ for all $x \in \mathcal{X}$ suffices for decision-making in contextual bandits. Most HTE estimators are designed to estimate a function $\hat{\tau}$ such that they minimize $\mathbb{E}_x[(\hat{\tau}(x) - \tau^*(x))^2]$. As one may expect, at any context x , the error $(\hat{\tau}(x) - \tau^*(x))^2$ in estimating treatment effects is likely to get large as $\min(p(a_0|x), p(a_1|x))$ gets closer to zero (i.e., less experimentation at context x). For example, the weights used by DR-learner (Kennedy, 2020) grow as action sampling probabilities shrink, leading to higher variance for this estimator. This implies that most HTE estimators suffer slow estimation rates for the contextual bandit setting, where action sampling probabilities need to shrink in order to obtain optimal cumulative regret guarantees.

We argue that our analysis of the R -learner allows us to side step this issue (when assignment probabilities are known) without assuming that assignment probabilities are bounded away from zero. In particular,

⁵Since excess risk bounds for empirical risk minimizers can be bounded by the complexity of the model class, Assumption 1 can be proved for several classes using results in Koltchinskii (2011).

one can easily show from Proposition 1 that the R -learner excess risk in the two arm setting is given by $\mathbb{E}_x[p(a_0|x)p(a_1|x)(\hat{\tau}(x) - \tau^*(x))^2]$. Unlike the previous estimand, note that $p(a_0|x)p(a_1|x)(\hat{\tau}(x) - \tau^*(x))^2$ goes to zero as the action sampling probability for either treatment or control goes to zero. Therefore, the R -loss excess risk can be bounded by fast rates even with arbitrarily small action-sampling probabilities (when they are known). This insight extends to the general multi-arm setting.

Ultimately, we show the R -loss excess risk is exactly the quantity we need to bound in order to obtain our optimal regret guarantees (see Lemma 5). Therefore, the additional properties we described in this section, of fast rates for the R -loss excess risk under vanishing action-sampling probabilities and tighter control on model misspecification error (stated in Proposition 2), lead us to identify the R -learner as an ideal HTE estimator for the contextual bandit setting.

4 ALGORITHM

Algorithm 1: Heterogeneous Treatment Effect Inverse Gap Weighting

Input: epoch schedule τ_m for all $m \in \mathbb{N}$, confidence parameter $\delta \in (0, 1)$

Initialize $\gamma_1 = 1$, $\hat{g}_1 \equiv 0$, $\tau_0 = 0$, $H_0 = \emptyset$, and $safe = \text{True}$.

for epoch $m = 1, 2, \dots$ **do**

for round $t = \tau_{m-1} + 1, \dots, \tau_m$ **do**

 Observe context x_t .

if $safe$ **then**

 Let kernel p_m be defined by (11) given γ_m and \hat{g}_m .

 Sample $a_t \sim p_m(\cdot|x_t)$ and observe $r_t(a_t)$.

 Update $H_t = H_{t-1} \cup \{(x_t, a_t, r_t(a_t))\}$.

 Update

$(safe, \hat{m}) \leftarrow \text{CheckAndChooseSafe}(H_t)$.

else

 Sample $a_t \sim p_{\hat{m}}(\cdot|x_t)$ and observe $r_t(a_t)$.

if $safe$ **then**

 Let

$\gamma_{m+1} = c\sqrt{K/\xi(\tau_m - \tau_{m-1}, \delta/(m+1)^2)}$.

 Let \hat{g}_{m+1} be the R -loss oracle output with input data $H_{\tau_m} \setminus H_{\tau_{m-1}}$ from epoch m .

In this section, we outline the Heterogeneous Treatment Effect Inverse Gap Weighting (HTE-IGW) algorithm for contextual bandits. A formal description is stated in Algorithm 1 below. Our algorithm is based on the FALCON+ algorithm in (Simchi-Levi and Xu, 2021, Algorithm 2) and the Safe-FALCON algorithm in (Krishnamurthy et al., 2021, Algorithm 1).

Epoch Schedule HTE-IGW is implemented using a schedule of epochs where epoch $m \in \mathbb{N}$ starts at round $\tau_{m-1} + 1$ and ends at round τ_m . This is done to reduce the number of calls to the R -loss oracle. For example, if we set $\tau_m = 2^m$, then the algorithm runs in $\mathcal{O}(\log T)$ epochs for any (possibly unknown) T . The algorithm accepts any valid epoch schedule. For any round $t \in \mathbb{N}$, we let $m(t)$ denote the epoch corresponding to round t .

Adaptive Behavior Each round t starts with a status that is called “safe” or “not safe”, depending on whether the algorithm has detected evidence of model misspecification. At the beginning of any epoch m that starts on a “safe” round, the algorithm uses an estimate of the treatment effect model \hat{g}_m , obtained via an offline R -loss oracle using data from the previous epoch $m - 1$. As long as the “safe” status is maintained, the algorithm assigns actions by drawing from the following *action selection kernel*:

$$p_m(a|x) := \begin{cases} \frac{1}{K + \gamma_m(\hat{g}_m(x, \hat{a}) - \hat{g}_m(x, a))} & \text{for } a \neq \hat{a} \\ 1 - \sum_{a' \neq \hat{a}} p(a'|x) & \text{for } a = \hat{a}, \end{cases} \quad (11)$$

where $\hat{a} = \arg \max_a \hat{g}_m(a, x)$ is the current predicted best action. The parameter $\gamma_m > 0$ governs how much the algorithm exploits and explores: assignment probabilities concentrate on the predicted best policy \hat{a} when γ_m is large and are more spread out when γ_m is small. During “safe” epochs, this parameter is set to be

$$\gamma_m := c\sqrt{K/\xi(\tau_{m-1} - \tau_{m-2}, \delta'/m^2)}, \quad (12)$$

where $c = \sqrt{1/8}$ is a fixed constant, ξ is the estimation rate of the R -loss oracle defined in (8), $\delta' = \delta/2$ where $\delta > 0$ is a confidence parameter, and the quantity $\tau_{m-1} - \tau_{m-2}$ is the size of the previous batch which was used to estimate the model \hat{g}_m . Definition (12) implies that small classes such as linear models allow for a quickly increasing γ_m and therefore more exploitation, while large classes require more exploration since γ_m increases more slowly.

If the algorithm ever switches to “not safe” status, it switches to a safe kernel that was previously used and never returns to “safe” status. This safe kernel guarantees a bounded level of regret. To determine the safety status of the bandit, we use a misspecification test.

Misspecification Test We make use of the same misspecification testing procedure developed in Krishnamurthy et al. (2021) to determine the safety status of every round. This misspecification test constructs a Hoeffding-style high-probability lower bound on the cumulative reward and triggers the safety check mechanism when the cumulative reward dips below this lower bound. If triggered, the safety mechanism defaults to the most recent kernel with the largest lower bound. Krishnamurthy et al. (2021) show that the safety test is triggered only after the

lower bounds on rewards are sufficiently large. Therefore, this high-probability safety test ensures that the expected regret stays bounded as long as the safety test is not triggered.

In our algorithm, this misspecification test is implemented as the CheckAndChooseSafe method which takes in the interaction history to return an indicator variable *safe*, indicating whether the misspecification test passed, and the safety epoch \hat{m} , pointing to the previously used safe kernel $p_{\hat{m}}$. Once the algorithm switches to “not safe” status, it continues to use the safe kernel and never returns to “safe” status. The details of an efficient implementation of this misspecification test can be viewed in Algorithms 2 and 3 of Krishnamurthy et al. (2021). For our purposes, it suffices to know that this test guarantees bounds on the instantaneous regret after the misspecification error surpasses the estimation error. Although the results are important, the details of the implementation are somewhat orthogonal to the emphasis of this paper. See Appendix B.5 for further discussion. Lastly, note that the misspecification level need not be known, though our regret bounds stated in Section 5 will depend on it.

5 MAIN RESULTS

The following provides a regret guarantee on our algorithm that depends on the known estimation rates and the misspecification error of the R -loss oracle.

Theorem 1. *Suppose the R -loss oracle satisfies Assumption 1. Then, with probability at least $1 - \delta$, HTE-IGW attains the following regret guarantee:*

$$\text{Reg}(T) \leq \mathcal{O} \left(\sum_{t=\tau_1+1}^T \sqrt{K\xi \left(\tau_{m(t)-1} - \tau_{m(t)-2}, \frac{\delta'}{m(t)^2} \right)} + \sqrt{KBT} \right). \quad (13)$$

Under epoch schedule $\tau_m = 2^m$, Theorem 1 gives

$$\text{Reg}(T) \leq \mathcal{O}(\sqrt{K\xi(T, \delta'/\log T)}T + \sqrt{KBT}). \quad (14)$$

The first term is sublinear in T (typically $\tilde{\mathcal{O}}(\sqrt{T})$ with problem-dependent parametric rates on excess risk) since the excess risk rate decreases in the sample size, and the second term is evidently linear in T . In the Appendix A.3, we discuss excess risk rates for the R -learner oracle under various important function classes. For a more concrete example, consider the following guarantees on the class of linear treatment effect models.

Example 1 (Linear treatment effect). Consider the following class of linear treatment effects

$$\mathcal{G} = \{g : (x, a) \mapsto \langle \theta, x_a \rangle \mid \theta \in \mathbb{R}^d, \|\theta\|_2 \leq 1\}$$

where $x = (x_a)_{a \in \mathcal{A}}$, $x_a \in \mathbb{R}^d$, and $\|x_a\|_2 \leq 1$. Under a linear class, the R -learner reduces to linear least squares regression of residual outcomes on residual treatments. Therefore, we can use the excess risk rates for square error minimization used in Simchi-Levi and Xu (2021) to get the same optimal rate guarantees under realizability while obtaining stronger misspecification bounds compared to Krishnamurthy et al. (2021) due to our results in Proposition 2. In particular, we have $\text{Reg}(T) \leq O(\sqrt{KT(d + \log T)} + \sqrt{KBT})$.

Note that Theorem 1 provides a bias-variance trade-off for contextual bandits. The first term in (13) is the regret bound under realizability and it depends on the estimation rate, which captures the variance of the R -loss oracle estimate over the class \mathcal{G} . The second term in (13) is the regret overhead due to misspecification and it depends on B , which captures a notion of bias for the best estimator in the model class \mathcal{G} under the distribution induced by any probability kernel. For more expressive model classes, the bias term B is small, but the variance term ξ is large, showing that the bounds quantify a bias-variance trade-off for contextual bandits that rely on heterogeneous treatment effect estimation. Moreover, note that Theorem 1 gives a tighter quantification of the bias-variance trade-off for contextual bandits compared to Krishnamurthy et al. (2021) since misspecification under the R -learner loss is no greater than misspecification under the squared error loss as shown by Proposition 2.

Limitations Our regret bound for the linear model class has worse dependence on the number of arms than the best known upper bound on stochastic contextual bandits (Li et al., 2019). This is an inherent limitation of the action-sampling strategy of inverse gap weighting Abe and Long (1999) which does not make use of the additional structure of the model class. Another limitation is that the exploration rate of our algorithm relies on knowledge of R -loss excess risk estimation rates for the specified model class. We discuss instance-dependent estimation rates of the R -learner under broad classes of models. However, for any estimator and any model class, fast estimation rates may not be known. Further work on R -loss excess risk fast rates or setting the exploration rate in a more data-driven manner would bring practical benefits to our current approach.

6 EXTENSIONS

Model Selection Our algorithm depends on specifying a model class with a known estimation rate. We can consider procedures for automatically selecting the appropriate model class in an adaptive manner. In particular, we could easily adapt many methods of model selection (Cutkosky et al., 2021; Foster et al., 2019; Krishnamurthy and Athey, 2021; Ghosh et al., 2021) to our setting. In our second ex-

periment in Section 7, we conduct a form of model selection on linear classes using LASSO regularization on the empirical R -loss risk objective. Using this, we adaptively set the excess risk rates based on the estimated sparsity (i.e., complexity) of the model.

Non-stationary Confounder Suppose the underlying action-independent component of the true conditional mean reward varied with time: $f_t^*(x, a) = g^*(x, a) + h_t^*(x)$ where h_t^* is a time-varying function of contexts. This scenario can occur in practice where there are non-stationary effects that are independent of treatment recommendations to experimental units such as time-of-day effects on digital platform traffic. Non-stationary confounding factors can make reward model estimation more difficult and thus affect bandit performance as noted in Qin and Russo (2022). Our algorithm is more robust to such scenarios since it is able to estimate a stationary treatment effect when it exists.

7 EXPERIMENTS

In our experiments, we compare the performance of the HTE-IGW algorithm with the similar inverse gap weighting algorithm (referred to as IGW here) that uses an offline least squares regression oracle to model the conditional reward function (Simchi-Levi and Xu, 2021; Krishnamurthy et al., 2021). The reason is to isolate the benefits of using treatment effect models to make fair comparisons.

7.1 Setup

We describe the common setup across experiments. Define $\mathcal{S}^d := \{x \in \mathbb{R}^d : \|x\|_2 = 1\}$ with $d = 100$. For simplicity, we set $\mathcal{X} = \mathcal{S}^d$ and $\mathcal{A} = \{1, 2\}$. At each time step $t \leq T$ for $T = 10,000$, the learner observes an independently sampled context $x_t \sim \text{Uniform}(\mathcal{X})$, takes an action $a_t \in \mathcal{A}$, and observes a reward $r_t(a_t) = f^*(x_t, a_t) + \varepsilon_t = g^*(x_t, a_t) + h^*(x_t) + \varepsilon_t$ where $\varepsilon_t \sim \sqrt{12\sigma^2} \times \text{Uniform}([-1/2, 1/2])$ independently with $\sigma = 1/10$. The specification of $f^* = g^* + h^*$ is provided in each scenario below. We estimate the conditional mean realized reward nuisance parameter $\hat{\mu}$ using cross-fitting. Notably, in every single experiment, we specify the linear function class $\mathcal{G} = \{g : (x, a) \mapsto \langle \theta, x_a \rangle \mid \theta \in \mathbb{R}^d\}$ for model estimation. We plot the total regret curve for each experiment. See the figures below. The shaded areas represent a standard deviation from the mean curve over 25 seed runs for each experiment.

7.2 Linear Reward & Linear Treatment Effect

First, we consider the setting where the conditional mean reward is linear in the contexts. This covers the baseline case where the reward and treatment effect have the same complexity and they are well-specified. We set $f^*(x, a) =$

$1 + \langle \theta_a, x \rangle$ where $\theta_a \sim \text{Uniform}(\mathcal{S}^d)$ independently for all $a \in \mathcal{A}$. In this case, $h^* \equiv 0$.

We find that both algorithms have almost exactly the same performance. Both algorithms exhibit similar sublinear total regret. This validates the observation that when the treatment effect and reward are equivalent, HTE-IGW performs just as well as a reward-based method under realizability. See Figure 1.




Figure 1: Total regret of synthetic experiment with linear reward and linear treatment effect.

7.3 Linear Reward & Constant Treatment Effect

Next, we consider the setting where the treatment effect model lies in a simpler model class than the reward model, yet model realizability still holds for both the treatment effect and reward models. Specifically, the treatment effects are constant and independent of context with a large gap between the treatments, and the confounding term is linear in contexts. Clearly, a linear model class is capable of capturing both the constant treatment effect and linear reward model. We set $g^*(x, a) = u_a + \mathbb{I}(a = 1)$ where $u_a \sim \text{Uniform}([0, 1])$ independently for all $a \in \mathcal{A}$ and $h^*(x) = 1 + \langle \theta, x \rangle$ where $\theta \sim \text{Uniform}(\mathcal{S}^d)$.

We find again that HTE-IGW performs no worse than IGW. In this case, the treatment effect model estimation procedure using a linear class is not capturing the sparsity of the true function. Therefore, we consider model selection to enforce simpler models and readjust exploration parameters according to the model complexity. See top figure in Figure 2.

Model Selection We consider the same problem setup, but we use LASSO regularization as well on the risk objective. We also set the exploration rates based on estimated sparsity of the model. Since there is a form of model selection being used here, we refer to the respective algorithms as MOD-HTE-IGW and MOD-IGW to make a distinction on their additional adaptivity beyond the base version of these algorithms.

In this case, we find that MOD-HTE-IGW significantly out-

performs MOD-IGW. The R -loss oracle quickly finds that none of the context features are predictive of the treatment effects and zeros out all the coefficients besides the bias in the model. Since the treatment effect gap is large, MOD-HTE-IGW uses a simple model to very quickly find the optimal action and thereby induces a less aggressive exploration parameter. On the other hand, the least squares oracle finds that all the context features are predictive of the rewards since the confounding term has non-zero coefficients for each context. This induces a more aggressive exploration parameter to ensure the algorithm learns an accurate model of a more complex reward function. Essentially, MOD-IGW over-explores on a less accurate model over longer period of time. We find that although both MOD-IGW and MOD-HTE-IGW have sublinear total regret in this example, there is a large gap between them with MOD-HTE-IGW showing significant improvements in performance. This validates that when the treatment effect model lies in a simpler class than the reward model, treatment effect-based methods (with model selection) can vastly outperform the reward-based method. See bottom figure in Figure 2.

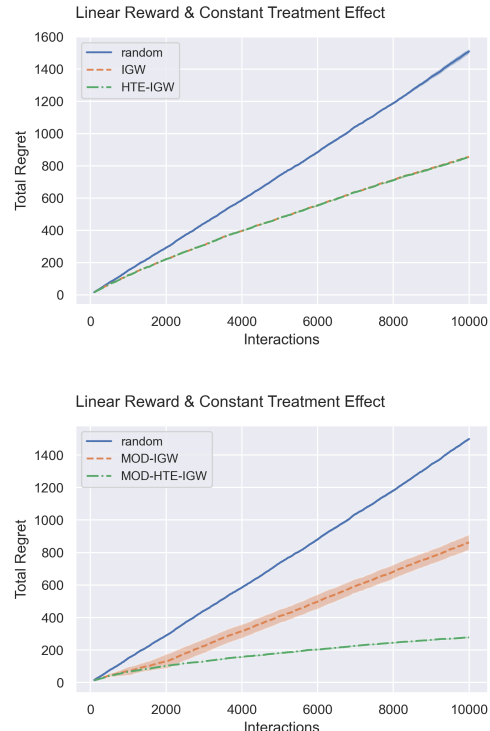


Figure 2: Total regret of synthetic experiment with linear reward and constant treatment effect. Top: w/o model selection. Bottom: w/ model selection.

7.4 Step-wise Linear Reward & Linear Treatment Effect

Lastly, we consider the setting where the treatment effect model lies in a simpler model class than the reward model, and the reward model is misspecified. Specifically, the treatment effects are linear and the confounding term is step-wise linear. The linear model class is capable of capturing the linear treatment effect model but not the step-wise linear reward model. We set $g^*(x, a) = u_a + \mathbb{I}(a = 1)$ where $u_a \sim \text{Uniform}([0, 1])$ independently for all $a \in \mathcal{A}$, and $h^*(x) = -\mathbb{I}(x_0 > 1/4) \cdot \max_a(1 + \langle \theta_a, x \rangle)$ where $\theta_a \sim \text{Uniform}(\mathcal{S}^d)$.

We find that due to misspecification of the reward model, the total regret of IGW is not much better than random sampling, but HTE-IGW manages to make noticeable improvements. See Figure 3.




Figure 3: Total regret of synthetic experiment with step-wise linear reward and linear treatment effect.

7.5 Perturbed Step-wise Linear Reward & Perturbed Linear Treatment Effect

We present an additional experiment for the case where the treatment effect model is also misspecified. We use a linear model class while the true treatment effects are linear with a sinusoidal perturbation and the confounding term is step-wise linear with a sinusoidal perturbation. Therefore, the true treatment effect and reward models are non-linear and misspecified under the linear model class. We set $g^*(x, a) = \mathbb{I}(a = 1) + \langle \theta_a, x \rangle + \sin(\langle \theta_a, x \rangle)$ and $h^*(x) = -\mathbb{I}(x_0 > 1/4) \cdot \max_a(1 + \langle \theta_a, x \rangle)$ where $\theta_a \sim \text{Uniform}(\mathcal{S}^d)$ independently for all $a \in \mathcal{A}$. We use a linear model class (with and without model selection). We observe that the performance of both HTE-IGW and IGW suffer from model misspecification. However, with model selection, HTE-IGW regains the performance, showing that HTE-IGW is more capable of capturing the simplest model class necessary for decision making even under model misspecification. See Figure 4.

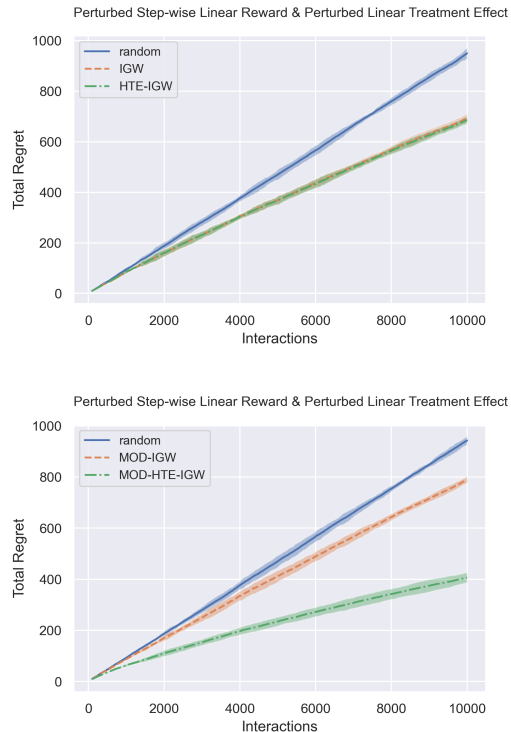


Figure 4: Total regret of synthetic experiment with non-linear (sinusoidally perturbed step-wise linear) reward and linear treatment effect. Top: w/o model selection. Bottom: w/ model selection.

8 CONCLUSION

In this paper, we presented the first universal reduction of cumulative regret minimization in stochastic contextual bandits to heterogeneous treatment effect estimation method via R -loss minimization. Based on this reduction, we develop a computationally efficient, flexible, general-purpose algorithm that leverages techniques presented in (Krishnamurthy et al., 2021) to not require treatment effect realizability and to be more robust to misspecification. Furthermore, we showed that treatment effect estimation oracles are a more robust primitive for contextual bandits than squared error regression oracles as commonly used in reward estimation (Foster and Rakhlin, 2020; Simchi-Levi and Xu, 2021). Beyond cumulative regret minimization, heterogeneous treatment effect estimation should also lead to better contextual bandit algorithms for simple regret minimization—which still appears to be a challenging problem based on existing algorithms (Athey et al., 2022). We hope our findings will inspire further research on adapting econometric methods of treatment effect estimation for developing more robust and practical contextual bandits.

Acknowledgements

The authors would like to thank Undral Byambadalai, Stefan Wager, Ruohan Zhan, and the reviewers for their helpful feedback. The authors are also grateful for the generous support provided by Golub Capital Social Impact Lab. S.A. and S.K.K. also acknowledge generous support from the Office of Naval Research grant N00014-19-1-2468.

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.
- Naoki Abe and Philip M Long. Associative reinforcement learning using linear probabilistic concepts. In *ICML*, pages 3–11. Citeseer, 1999.
- Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, pages 1638–1646. PMLR, 2014.
- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International conference on machine learning*, pages 127–135. PMLR, 2013.
- Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47(2): 1148–1178, 2019.
- Susan Athey, Undral Byambadalai, Vitor Hadad, Sanath Kumar Krishnamurthy, Weiwen Leung, and Joseph Jay Williams. Contextual bandits in a survey experiment on charitable giving: Within-experiment outcomes versus policy learning. *arXiv preprint arXiv:2211.12004*, 2022.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- Peter L Bartlett, Olivier Bousquet, and Shahar Mendelson. Local rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.
- Keith Battocchi, Eleanor Dillon, Maggie Hei, Greg Lewis, Paul Oka, Miruna Oprescu, and Vasilis Syrgkanis. Econml: A python package for ml-based heterogeneous treatment effects estimation. <https://github.com/microsoft/EconML>, 2019. Version 0.x.
- Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 19–26. JMLR Workshop and Conference Proceedings, 2011.
- Alberto Bietti, Alekh Agarwal, and John Langford. A contextual bandit bake-off. *Journal of Machine Learning Research*, 22(133):1–49, 2021.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James M. Robins. Double/debiased machine learning for treatment and structural parameters. *Econometrics: Econometric & Statistical Methods - General eJournal*, 2017.
- Young-Geun Choi, Gi-Soo Kim, Seunghoon Paik, and Myunghee Cho Paik. Semi-parametric contextual bandits with graph-laplacian regularization. *arXiv preprint arXiv:2205.08295*, 2022.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings, 2011.
- Coba. <https://github.com/VowpalWabbit/coba>, 2022.
- Ashok Cutkosky, Christoph Dann, Abhimanyu Das, Claudio Gentile, Aldo Pacchiano, and Manish Purohit. Dynamic balancing for model selection in bandits and rl. In *International Conference on Machine Learning*, pages 2276–2285. PMLR, 2021.
- Miroslav Dudík, Daniel Hsu, Satyen Kale, Nikos Karampatziakis, John Langford, Lev Reyzin, and Tong Zhang. Efficient optimal learning for contextual bandits. *arXiv preprint arXiv:1106.2369*, 2011.
- Miroslav Dudík, Katja Hofmann, Robert E Schapire, Aleksandrs Slivkins, and Masrour Zoghi. Contextual dueling bandits. In *Conference on Learning Theory*, pages 563–587. PMLR, 2015.
- Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. *Advances in Neural Information Processing Systems*, 23, 2010.
- Dylan Foster and Alexander Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 3199–3210. PMLR, 2020.
- Dylan Foster, Alekh Agarwal, Miroslav Dudík, Haipeng Luo, and Robert Schapire. Practical contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 1539–1548. PMLR, 2018.
- Dylan J Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*, 2019.
- Dylan J Foster, Akshay Krishnamurthy, and Haipeng Luo. Model selection for contextual bandits. *Advances in Neural Information Processing Systems*, 32, 2019.

- Dylan J Foster, Claudio Gentile, Mehryar Mohri, and Julian Zimmert. Adapting to misspecification in contextual bandits. *Advances in Neural Information Processing Systems*, 33:11478–11489, 2020.
- Avishek Ghosh, Sayak Ray Chowdhury, and Aditya Gopalan. Misspecified linear bandits. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Avishek Ghosh, Abishek Sankararaman, and Ramchandran Kannan. Problem-complexity adaptive model selection for stochastic linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 1396–1404. PMLR, 2021.
- Kristjan Greenewald, Ambuj Tewari, Susan Murphy, and Predag Klasnja. Action centered contextual bandits. *Advances in neural information processing systems*, 30, 2017.
- Yonatan Gur, Ahmadreza Momeni, and Stefan Wager. Smoothness-adaptive contextual bandits. *Operations Research*, 2022.
- Yichun Hu, Nathan Kallus, and Xiaojie Mao. Smooth contextual bandits: Bridging the parametric and non-differentiable regret regimes. In *Conference on Learning Theory*, pages 2007–2010. PMLR, 2020.
- Edward H Kennedy. Optimal doubly robust estimation of heterogeneous causal effects. *arXiv preprint arXiv:2004.14497*, 2020.
- Vladimir Koltchinskii. *Oracle inequalities in empirical risk minimization and sparse recovery problems: École D’Été de Probabilités de Saint-Flour XXXVIII-2008*, volume 2033. Springer Science & Business Media, 2011.
- Vladimir Koltchinskii and Dmitriy Panchenko. Rademacher processes and bounding the risk of function learning. In *High dimensional probability II*, pages 443–457. Springer, 2000.
- Akshay Krishnamurthy, Zhiwei Steven Wu, and Vasilis Syrgkanis. Semiparametric contextual bandits. In *International Conference on Machine Learning*, pages 2776–2785. PMLR, 2018.
- Sanath Kumar Krishnamurthy and Susan Athey. Optimal model selection in contextual bandits with many classes via offline oracles. *arXiv preprint arXiv:2106.06483*, 2021.
- Sanath Kumar Krishnamurthy, Vitor Hadad, and Susan Athey. Adapting to misspecification in contextual bandits with offline regression oracles. In *International Conference on Machine Learning*, pages 5805–5814. PMLR, 2021.
- Sören R Künnel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10):4156–4165, 2019.
- John Langford and Tong Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. *Advances in neural information processing systems*, 20, 2007.
- Tor Lattimore, Csaba Szepesvari, and Gellert Weisz. Learning with good feature representations in bandits and in rl with a generative model. In *International Conference on Machine Learning*, pages 5662–5670. PMLR, 2020.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.
- Yingkai Li, Yining Wang, and Yuan Zhou. Nearly minimax-optimal regret for linearly parameterized bandits. In *Conference on Learning Theory*, pages 2173–2174. PMLR, 2019.
- Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2): 299–319, 2021.
- Yi Peng, Miao Xie, Jiahao Liu, Xuying Meng, Nan Li, Cheng Yang, Tao Yao, and Rong Jin. A practical semi-parametric contextual bandit. In *IJCAI*, pages 3246–3252, 2019.
- Chao Qin and Daniel Russo. Adaptivity and confounding in multi-armed bandit experiments. *arXiv preprint arXiv:2202.09036*, 2022.
- Peter M Robinson. Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, pages 931–954, 1988.
- Aadirupa Saha and Akshay Krishnamurthy. Efficient and optimal algorithms for contextual dueling bandits under realizability. In *International Conference on Algorithmic Learning Theory*, pages 968–994. PMLR, 2022.
- David Simchi-Levi and Yunzong Xu. Bypassing the monster: A faster and simpler optimal algorithm for contextual bandits under realizability. *Mathematics of Operations Research*, 2021.
- Yunbei Xu and Assaf Zeevi. Towards optimal problem dependent generalization error bounds in statistical learning theory. *arXiv preprint arXiv:2011.06186*, 2020.

A R-LOSS EXCESS RISK, MISSPECIFICATION, AND ESTIMATION RATES

A.1 Equivalent Form of Excess Risk

Proposition 1 states an equivalent form for the R -loss excess risk. Notice how this form is independent of the conditional mean realized reward estimate $\hat{\mu}$ that is fixed with respect to the data-generating distribution (as in a constant or cross-fitted estimate). Proposition 1 will later allow us to prove Proposition 2 and to bound in Lemma 5 the error of the treatment effect estimate using bounds on the excess risk of the R -loss oracle estimate.

Proposition 1. *Suppose $\hat{\mu}$ is any fixed function. For any $g^* \in \arg \min_{g'} \mathcal{R}_p(g')$,*

$$\mathcal{E}_p(g) = \mathbb{E} \left[\langle e_a - p(x), g(x, \cdot) - g^*(x, \cdot) \rangle^2 \right], \quad (6)$$

where the expectation is taken with respect to $D(p)$.

Proof. Consider any $g : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$. For conciseness in this proof, we shorten the notation $g(x) := g(x, \cdot) = (g(x, a))_{a \in \mathcal{A}}$. First of all, we have

$$\begin{aligned} \mathcal{R}_p(g) &= \mathbb{E} \left[(r(a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle)^2 \right] \\ &= \mathbb{E} \left[(r(a) - f^*(x, a) + f^*(x, a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle)^2 \right] \\ &= \mathbb{E} \left[(r(a) - f^*(x, a))^2 \right] + \mathbb{E} \left[(f^*(x, a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle)^2 \right] \\ &\quad - 2 \mathbb{E} \left[(r(a) - f^*(x, a)) \cdot (f^*(x, a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle) \right]. \end{aligned}$$

Since $\hat{\mu}$ is independent of $(x, a, r(a))$, the last term of this last equation is

$$\begin{aligned} &\mathbb{E} \left[(r(a) - f^*(x, a)) \cdot (f^*(x, a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle) \right] \\ &= \mathbb{E}_{x, a} \left[\mathbb{E}_r \left[(r(a) - f^*(x, a)) \cdot (f^*(x, a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle) \mid x, a \right] \right] \\ &= \mathbb{E}_{x, a} \left[(\mathbb{E}_r[r(a) \mid x, a] - f^*(x, a)) \cdot (f^*(x, a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle) \right] \\ &= \mathbb{E}_{x, a} \left[(f^*(x, a) - f^*(x, a)) \cdot (f^*(x, a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle) \right] \\ &= 0 \end{aligned}$$

and the second term is

$$\begin{aligned} &\mathbb{E} \left[(f^*(x, a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle)^2 \right] \\ &= \mathbb{E} \left[\left(\left(f^*(x, a) - \mathbb{E}_{a' \sim p(\cdot | x)} [f^*(x, a')] \right) - \langle e_a - p(x), g(x) \rangle + \left(\mathbb{E}_{a' \sim p(\cdot | x)} [f^*(x, a')] - \hat{\mu}(x) \right) \right)^2 \right] \\ &= \mathbb{E} \left[\left(\langle e_a - p(x), f^*(x) \rangle - \langle e_a - p(x), g(x) \rangle + (\mu(x) - \hat{\mu}(x)) \right)^2 \right] \\ &= \mathbb{E} \left[\left(\langle e_a - p(x), f^*(x) - g(x) \rangle + (\mu(x) - \hat{\mu}(x)) \right)^2 \right] \\ &= \mathbb{E} \left[\langle e_a - p(x), f^*(x) - g(x) \rangle^2 \right] + \mathbb{E} \left[(\mu(x) - \hat{\mu}(x))^2 \right] + 2 \mathbb{E} \left[(\mu(x) - \hat{\mu}(x)) \cdot \langle e_a - p(x), f^*(x) - g(x) \rangle \right] \\ &= \mathbb{E} \left[\langle e_a - p(x), f^*(x) - g(x) \rangle^2 \right] + \mathbb{E} \left[(\mu(x) - \hat{\mu}(x))^2 \right] + 2 \mathbb{E}_x \left[\mathbb{E}_{a \sim p(\cdot | x)} \left[(\mu(x) - \hat{\mu}(x)) \cdot \langle e_a - p(x), f^*(x) - g(x) \rangle \mid x \right] \right] \\ &= \mathbb{E} \left[\langle e_a - p(x), f^*(x) - g(x) \rangle^2 \right] + \mathbb{E} \left[(\mu(x) - \hat{\mu}(x))^2 \right] + 2 \mathbb{E}_x \left[(\mu(x) - \hat{\mu}(x)) \cdot \mathbb{E}_{a \sim p(\cdot | x)} \left[\langle e_a - p(x), f^*(x) - g(x) \rangle \mid x \right] \right] \\ &= \mathbb{E} \left[\langle e_a - p(x), f^*(x) - g(x) \rangle^2 \right] + \mathbb{E} \left[(\mu(x) - \hat{\mu}(x))^2 \right] + 2 \mathbb{E}_x \left[(\mu(x) - \hat{\mu}(x)) \cdot \langle p(x) - p(x), f^*(x) - g(x) \rangle \right] \\ &= \mathbb{E} \left[\langle e_a - p(x), f^*(x) - g(x) \rangle^2 \right] + \mathbb{E} \left[(\mu(x) - \hat{\mu}(x))^2 \right]. \end{aligned}$$

Therefore,

$$\mathcal{R}_p(g) = \mathbb{E} \left[(r(a) - f^*(x, a))^2 \right] + \mathbb{E} \left[\langle e_a - p(x), f^*(x) - g(x) \rangle^2 \right] + \mathbb{E} \left[(\mu(x) - \hat{\mu}(x))^2 \right].$$

This implies that the set of true treatment effect models $\arg \min_{g'} \mathcal{R}_p(g')$ are the functions g that satisfy $\langle e_a - p(x), g(x) \rangle = \langle e_a - p(x), f^*(x) \rangle$ for any $(x, a) \in \mathcal{X} \times \mathcal{A}$. This is the set of true treatment effect functions

$$\mathcal{G}^* = \{g : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1] \mid \exists h : \mathcal{X} \rightarrow [0, 1] \text{ s.t. } g(x, a) + h(x) = f^*(x, a), \forall (x, a) \in \mathcal{X} \times \mathcal{A}\}.$$

Therefore, for any $g^* \in \arg \min_{g'} \mathcal{R}_p(g')$,

$$\begin{aligned} \mathcal{E}_p(g) &= \mathcal{R}_p(g) - \mathcal{R}_p(g^*) \\ &= \mathbb{E}[\langle e_a - p(x), f^*(x) - g(x) \rangle^2] - \mathbb{E}[\langle e_a - p(x), f^*(x) - g^*(x) \rangle^2] \\ &= \mathbb{E}[\langle e_a - p(x), g^*(x) - g(x) \rangle^2]. \end{aligned}$$

□

Note that from the proof we can see that this form of the excess risk is equivalent to the excess risk that uses the true nuisance parameters inside the loss function. Since the rest of our analysis relies on bounds on the excess risk, this shows that we lose nothing by relying on this form of the R -loss risk in our analysis, assuming the estimate $\hat{\mu}$ of μ is fixed.

A.2 Bound on Misspecification Error

Proposition 2 upper bounds the excess risk of the R -learner loss by the excess risk of the squared error loss.

Proposition 2. *For a fixed model class \mathcal{G} , the average misspecification error under the R -loss is no greater than the average misspecification error under the squared error loss, i.e.,*

$$B \leq \max_{p \in \mathcal{P}} \min_{f \in \mathcal{G}} \mathbb{E} \left[(f(x, a) - f^*(x, a))^2 \right]. \quad (7)$$

where the expectation is taken with respect to $D(p)$.

Proof. Let $g \in \mathcal{G}$ and $g^* \in \arg \min_{g'} \mathcal{R}_p(g')$. Also, let $\Delta_{x,a} := g(x, a) - g^*(x, a)$. Then, from Proposition 1 we have

$$\begin{aligned} \mathcal{E}_p(g) &= \mathbb{E} [\langle e_a - p(x), g(x, \cdot) - g^*(x, \cdot) \rangle^2] \\ &= \mathbb{E} \left[\left((g(x, a) - g^*(x, a)) - \mathbb{E}_{a' \sim p(\cdot|x)} [g(x, a') - g^*(x, a')] \right)^2 \right] \\ &= \mathbb{E} \left[\left(\Delta_{x,a} - \mathbb{E}_{a' \sim p(\cdot|x)} [\Delta_{x,a'}] \right)^2 \right] \\ &= \mathbb{E} \left[\Delta_{x,a}^2 + \mathbb{E}_{a' \sim p(\cdot|x)} [\Delta_{x,a'}]^2 - 2 \mathbb{E}_{a' \sim p(\cdot|x)} [\Delta_{x,a'}] \cdot \Delta_{x,a} \right] \\ &= \mathbb{E} [\Delta_{x,a}^2] + \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\mathbb{E}_{a \sim p(\cdot|x)} [\Delta_{x,a}]^2 \right] - 2 \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\mathbb{E}_{a' \sim p(\cdot|x)} [\Delta_{x,a'}] \cdot \mathbb{E}_{a \sim p(\cdot|x)} [\Delta_{x,a} | x] \right] \\ &= \mathbb{E} [\Delta_{x,a}^2] + \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\mathbb{E}_{a \sim p(\cdot|x)} [\Delta_{x,a}]^2 \right] - 2 \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\mathbb{E}_{a \sim p(\cdot|x)} [\Delta_{x,a}]^2 \right] \\ &= \mathbb{E} [\Delta_{x,a}^2] - \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\mathbb{E}_{a \sim p(\cdot|x)} [\Delta_{x,a}]^2 \right] \\ &= \mathbb{E} \left[(g(x, a) - g^*(x, a))^2 \right] - \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\mathbb{E}_{a \sim p(\cdot|x)} [g(x, a) - g^*(x, a)]^2 \right] \\ &\leq \mathbb{E} \left[(g(x, a) - g^*(x, a))^2 \right]. \end{aligned}$$

Since $f^* \in \arg \min_{g'} \mathcal{R}_p(g')$, this implies that

$$B \leq \max_{p \in \mathcal{P}} \min_{f \in \mathcal{G}} \mathbb{E} \left[(f(x, a) - f^*(x, a))^2 \right].$$

□

In other words, for a chosen model class \mathcal{G} , the average misspecification error of the R -loss risk is upper bounded by the average misspecification error of the square error risk. A square error oracle is usually used to estimate the reward model in Simchi-Levi and Xu (2021). This implies that heterogeneous treatment effect estimation via R -loss oracles is more robust to model misspecification than reward estimation via least squares oracles. This can be interpreted to mean that heterogeneous treatment effect estimation is no harder than reward estimation.

Misspecification Error Bound Intuition We provide further intuition to the previous result. Essentially, the set of true treatment effects

$$\mathcal{G}^* := \arg \min_{g'} \mathcal{R}_p(g') = \{g : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1] \mid \exists h : \mathcal{X} \rightarrow [0, 1] \text{ s.t. } g(x, a) + h(x) = f^*(x, a), \forall (x, a) \in \mathcal{X} \times \mathcal{A}\}$$

specifies a buffer zone around the true reward model f^* of sufficiently valid estimands for decision-making. This set contains infinitely many possible target estimands, including the true reward model f^* . Therefore, the “distance” B_R from any fixed model class \mathcal{G} to the “nearest” estimand g^* (corresponding to the misspecification error under the R -loss, i.e., the left hand side of inequality (7)) in this buffer zone must necessarily be no further than the “distance” B_{sq} to the true reward model f^* (corresponding to the misspecification error under the squared error loss, i.e., the right-hand side of inequality (7)). See Figure 5 for a graphical representation of this observation. Note that \mathcal{G} may possibly not contain f^* even while at the same time possibly still overlap with \mathcal{G}^* , in which case the treatment effect model would be well-specified while the reward model would be misspecified. If the reward model is well-specified, then the treatment effect model is well-specified. However, the converse may not necessarily be true.




Figure 5: Intuition behind differences between R -loss vs. squared error loss misspecification error. B_R denotes the misspecification error of the R -loss, and B_{sq} denotes the misspecification error of the squared error loss.

A.3 R-learner Oracle Estimation Rates

Bartlett et al. (2005) and Xu and Zeevi (2020) discuss problem-dependent rates on the excess risk of a model trained via empirical risk minimization. We leverage their loss-dependent rates (see Equation 4.3 in Xu and Zeevi, 2020) and upper bound them with probability $1 - \delta$ as

$$\mathcal{E}_p(\hat{g}) \leq \mathcal{O}\left(r^* \vee \sqrt{B r^*}\right) \leq \mathcal{O}(r^* + B) \quad (15)$$

where B is the average misspecification error under the given loss and r^* is the fixed point of a “surrogate” function $\psi(r; \delta)$ that upper bounds the uniform error within a localized region of the model class. Note that we allow for a loose bound

on the dependence on B because our regret guarantees are stated with an additive regret overhead due to misspecification that depends linearly on B , so this error can be absorbed into the misspecification overhead. This also makes it clear how misspecification issues can be separated from estimation in our algorithm since the bound on regret due to misspecification is covered by the safety check guarantees and so misspecification does not need to shape the exploration strategy to achieve our regret bounds.

Clearly, the fixed point depends on the model and the probability level δ . However, note that the following bounds on the fixed point of the surrogate function hide a logarithmic dependence on δ , typically of the form $\log(1/\delta)$ for the sake of simplicity as well. For VC-type classes defined by the class of models whose metric entropy is $\mathcal{O}(d \log \frac{1}{\epsilon})$ for any covering net size $\epsilon > 0$ where d is the VC dimension, it can be proven (Koltchinskii and Panchenko, 2000; Xu and Zeevi, 2020) that $r^* \leq \mathcal{O}(\frac{d \log n}{n})$. Therefore, VC-type classes have R -loss excess risk bound $\mathcal{O}(\frac{d \log n}{n} + B)$. Note that linear model classes with d -dimensional coefficients are VC-type classes with VC dimension equal to d . Similarly, for any general non-parametric class of models whose metric entropy is $\mathcal{O}(\epsilon^{-2\rho})$ for any covering net size $\epsilon > 0$ where $\rho \in (0, 1)$ is a constant, it can be proven that $r^* \leq \mathcal{O}(n^{-\frac{1}{1+\rho}})$. Therefore, non-parametric classes of polynomial growth have R -loss excess risk bound $\mathcal{O}(n^{-\frac{1}{1+\rho}} + B)$.

B ALGORITHM ANALYSIS

B.1 Definitions

Recall we have a class of treatment effect models $\mathcal{G} \subset \{g : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]\}$ and the set of true treatment effect models is $\mathcal{G}^* := \arg \min_{g'} \mathcal{R}_p(g')$ where the infimum is taken with respect to all functions. The following analysis is stated with respect to a fixed $g^* \in \mathcal{G}^*$. This is without a loss of generalization because the subsequent results are the same independent of the choice of $g^* \in \mathcal{G}^*$. With an abuse of notation, for any function $g : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$, we let $g(x)$ denote the vector $g(x, \cdot) = (g(x, a))_{a \in \mathcal{A}}$.

A policy $\pi : \mathcal{X} \rightarrow \mathcal{A}$ is a deterministic mapping from contexts to actions. Let $\Psi = \mathcal{A}^{\mathcal{X}}$ be the universal policy space containing all possible policies. The policy π_g that is induced by a model $g \in \mathcal{G}$ is given by $\pi_g(x) := \arg \max_a g(x, a)$ for every $x \in \mathcal{X}$. Let π^* denote the policy induced by g^* , i.e., $\pi^* := \arg \max_a g^*(x, a)$. The expected instantaneous treatment effect of a policy $\pi \in \Psi$ with respect to a treatment effect model g and probability kernel $p : \mathcal{X} \rightarrow \Delta_{\mathcal{A}}$ is defined as

$$\Phi_p^g(\pi) := \mathbb{E}_{x \sim D_{\mathcal{X}}} [\langle e_{\pi(x)} - p(x), g(x) \rangle] = \mathbb{E}_{x \sim D_{\mathcal{X}}} [g(x, \pi(x)) - \sum_{a \in \mathcal{A}} p(a|x)g(x, a)]. \quad (16)$$

Moreover, we write $\Phi_p(\pi)$ to mean $\Phi_p^{g^*}(\pi)$, the true expected instantaneous treatment effect for policy π under probability kernel p . We define the expected instantaneous regret of a policy π with respect to the treatment effect model g as

$$\text{Reg}_g(\pi) := \Phi_p^g(\pi_g) - \Phi_p^g(\pi) = \mathbb{E}_{x \sim D_{\mathcal{X}}} [g(x, \pi_g(x)) - g(x, \pi(x))]. \quad (17)$$

Note that by this definition the expected instantaneous regret for any policy is independent of the choice of probability kernel. We write $\text{Reg}(\pi)$ to mean $\text{Reg}_{g^*}(\pi)$, the true expected instantaneous regret for policy π .

We also consider the true expected instantaneous reward of a policy π defined as

$$R(\pi) := \mathbb{E}_{x \sim D_{\mathcal{X}}} [f^*(x, \pi(x))]. \quad (18)$$

Note that, according to the discussion in Appendix A.1 on the set of true treatment effect models, the following equivalences for the true expected instantaneous regret hold:

$$\text{Reg}(\pi) = \Phi_p(\pi^*) - \Phi_p(\pi) = \mathbb{E}_{x \sim D_{\mathcal{X}}} [g^*(x, \pi^*(x)) - g^*(x, \pi(x))] \quad (19)$$

$$= R(\pi^*) - R(\pi) = \mathbb{E}_{x \sim D_{\mathcal{X}}} [f^*(x, \pi^*(x)) - f^*(x, \pi(x))]. \quad (20)$$

Remark 2. Although it is constructed differently through treatment effect models, this definition of the true expected instantaneous regret $\text{Reg}(\pi)$ in fact is equal to the true expected instantaneous regret in Simchi-Levi and Xu (2021); Krishnamurthy et al. (2021). This allows us to carry out more or less a similar analysis in our proofs.

Next, for any probability kernel p and any policy π , we let $V(p, \pi)$ denote the expected inverse probability of policy π under kernel p ,⁶ defined as

$$V(p, \pi) := \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\frac{1}{p(\pi(x)|x)} \right]. \quad (21)$$

This is a crucial quantity in our analysis that measures the overlap between a policy and a data-sampling kernel. Next, we let Γ_t denote the set of observations up to and including time t , that is,

$$\Gamma_t := \{(x_s, a_s, r_s(a_s))\}_{s=1}^t. \quad (22)$$

Lastly, we let $m(t)$ denote the epoch associated to round t .

B.2 Action Selection Kernel as a Randomized Policy

Given any probability kernel p , from Lemma 3 in Simchi-Levi and Xu (2021), there exists a unique product probability measure on Ψ , given by:

$$Q_p(\pi) := \prod_{x \in \mathcal{X}} p(\pi(x)|x). \quad (23)$$

This measure satisfies the following property

$$p(a|x) = \sum_{\pi \in \Psi} \mathbb{I}\{\pi(x) = a\} Q_p(\pi). \quad (24)$$

Since any probability kernel p induces the distribution Q_p over the set of deterministic policies Ψ , we can think of Q_p as a randomized policy induced by p . Equations (24) and (23) establish a correspondence between the probability kernel p and the induced randomized policy Q_p . For short hand, we let $Q_m \equiv Q_{p_m}$ for any epoch m .

B.3 Safe Epoch

We say an epoch m is safe when the status of the algorithm at the end of the epoch is safe, that is, when the variable *safe* is still set to True at the end of this epoch. Let \hat{m} be the last safe epoch determined by the CheckAndChooseSafe subroutine in Algorithm 1. Note that, for all $m \leq \hat{m}$, the epoch m is safe. Now let m^* be such that:

$$m^* := \max \left\{ m \mid C_0 B \leq \xi \left(\tau_m - \tau_{m-1}, \frac{\delta'}{m^2} \right) \right\}. \quad (25)$$

The epoch m^* is critical to our theoretical analysis. It is the last epoch before the estimation rate of the R -loss oracle goes below the misspecification level.

B.4 High Probability Event

Theorem 1 provides certain high probability bounds on cumulative regret. As a preliminary step in these proofs, it will be helpful to show that the event \mathcal{W}_1 defined below hold with high probability. Event \mathcal{W}_1 ensures that our treatment effect estimates are “good” models for the first few epochs. This describes the event where, for any epoch $m \leq \min(m^*, \hat{m})$, the R -loss risk of the estimated model \hat{g}_{m+1} can be bounded purely in terms of the known estimation rate of the algorithm:

$$\mathcal{W}_1 := \left\{ \forall m \leq \min(m^*, \hat{m}), \mathcal{E}_{p_m}(\hat{g}_{m+1}) \leq 2\xi \left(\tau_m - \tau_{m-1}, \frac{\delta'}{m^2} \right) \right\}. \quad (26)$$

Lemma 1. *Suppose the R -loss oracle satisfies Assumption 1. Then the event \mathcal{W}_1 holds with probability at least $1 - \delta$.*

Proof. Consider any epoch m such that epoch $m - 1$ was safe. Since epoch $m - 1$ was safe, for the first time-step in epoch m the algorithm samples an action from the probability kernel p_m . Hence from Assumption 1, with probability $1 - \delta'/m^2$, we have:

$$\mathcal{E}_{p_m}(\hat{g}_{m+1}) \leq \xi \left(\tau_m - \tau_{m-1}, \frac{\delta'}{m^2} \right) + C_0 B. \quad (27)$$

⁶In Simchi-Levi and Xu (2021), this term is called the decisional divergence between the randomized policy Q_p and deterministic policy π .

Therefore, the probability that (27) does not hold for some epoch $m \leq \min(m^*, \hat{m})$ can be bounded by

$$\sum_{m=1}^{\infty} \frac{\delta'}{m^2} \leq \delta.$$

Recall $\delta' = \delta/2$. Hence, from the definition of m^* , we get that \mathcal{W}_1 holds with probability at least $1 - \delta$. \square

B.5 Safety Check Assumptions

We will assume the CheckAndChooseSafe subroutine determines the status of the algorithm such that the following assumptions hold. See Algorithms 2 and 3 of Krishnamurthy et al. (2021) for an example of a safety check algorithm that guarantees these properties and see Lemmas 9 and 10 for their proofs.

Assumption 2. *Suppose the event \mathcal{W}_1 holds with probability at least $1 - \delta$. Then, for $t \leq \tau_{m^*+1}$, the status of the algorithm at the end of round t is safe, and when $t > \tau_{m^*+1}$, we have that $m^* + 1 \leq \hat{m}$ with probability at least $1 - \delta$.*

Assumption 3. *Suppose the event \mathcal{W}_1 holds with probability at least $1 - \delta$. Then, for any epoch $m \geq m^*$, the expected implicit regret is upper bounded as*

$$\sum_{\pi \in \Psi} Q_{p_m}(\pi) \text{Reg}(\pi) \leq O(\sqrt{KB}) \quad (28)$$

with probability at least $1 - \delta$.

Assumption 2 guarantees that the safe kernel will not be triggered before true safe epoch m^* . Assumption 3 guarantees that once the algorithm status is not safe, the expected instantaneous reward is at least controlled by the misspecification error.

Krishnamurthy et al. (2021) achieves the bound on the expected implicit regret for every epoch $m \geq m^*$ as stated in Assumption 3 by constructing a Hoeffding-style high-probability lower bound on the cumulative reward and triggering a safety check mechanism when the cumulative reward dips below this lower bound. Then, the safety mechanism defaults to the most recent kernel with the largest lower bound. This high-probability misspecification test ensures that the expected regret does not only stay bounded after the test is triggered, but also that the expected regret stays bounded before the test is triggered, even after the safe epoch has occurred.

B.6 Per-epoch Properties

Lemma 2 forms an equivalence of the expected instantaneous regret of a policy as the expected implicit regret of the randomized policy, which will be useful in bounding the total regret.

Lemma 2. *For any round t ,*

$$\mathbb{E}_{x_t, r_t, a_t} [r_t(\pi^*(x_t)) - r_t(a_t) | \Gamma_{t-1}] = \sum_{\pi \in \Psi} Q_{m(t)}(\pi) \text{Reg}(\pi).$$

Proof. Consider any time step $t \geq 1$, then from (24) relating kernel p and Q_p and (19) we have the following equalities:

$$\begin{aligned} & \mathbb{E}_{x_t, r_t, a_t} [r_t(\pi^*(x_t)) - r_t(a_t) | \Gamma_{t-1}] \\ &= \mathbb{E}_{x \sim D_{\mathcal{X}}, a \sim p_t(\cdot|x)} [f^*(x, \pi^*(x)) - f^*(x, a)] \\ &= \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\sum_{a \in \mathcal{A}} p_t(a|x) (f^*(x, \pi^*(x)) - f^*(x, a)) \right] \\ &= \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\sum_{a \in \mathcal{A}} \sum_{\pi \in \Psi} \mathbb{I}(\pi(x) = a) Q_{m(t)}(\pi) (f^*(x, \pi^*(x)) - f^*(x, a)) \right] \\ &= \sum_{\pi \in \Psi} Q_{m(t)}(\pi) \mathbb{E}_{x \sim D_{\mathcal{X}}} [f^*(x, \pi^*(x)) - f^*(x, \pi(x))] \\ &= \sum_{\pi \in \Psi} Q_{m(t)}(\pi) \text{Reg}(\pi). \end{aligned}$$

\square

Lemma 3 states a key bound on the estimated implicit regret of the randomized policy Q_m .

Lemma 3. For any epoch m ,

$$\sum_{\pi \in \Psi} Q_m(\pi) \text{Reg}_{\hat{g}_m}(\pi) \leq \frac{K}{\gamma_m}.$$

Proof. This follows essentially from unpacking the definitions of regret (17) and the representation of the action selection kernel p_m in the algorithm and (24):

$$\begin{aligned} & \sum_{\pi \in \Psi} Q_m(\pi) \text{Reg}_{\hat{g}_m}(\pi) \\ &= \sum_{\pi \in \Psi} Q_m(\pi) \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\hat{g}_m(x, \pi_{\hat{g}_m}(x)) - \hat{g}_m(x, \pi(x)) \right] \\ &= \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\sum_{\pi \in \Psi} Q_m(\pi) \left(\hat{g}_m(x, \pi_{\hat{g}_m}(x)) - \hat{g}_m(x, \pi(x)) \right) \right] \\ &= \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\sum_{a \in \mathcal{A}} \sum_{\pi \in \Psi} \mathbb{I}(\pi(x) = a) Q_m(\pi) \left(\hat{g}_m(x, \pi_{\hat{g}_m}(x)) - \hat{g}_m(x, a) \right) \right] \\ &= \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\sum_{a \in \mathcal{A}} p_m(a|x) \left(\hat{g}_m(x, \pi_{\hat{g}_m}(x)) - \hat{g}_m(x, a) \right) \right] \\ &= \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\sum_{a \in \mathcal{A}} \frac{\left(\hat{g}_m(x, \pi_{\hat{g}_m}(x)) - \hat{g}_m(x, a) \right)}{K + \gamma_m \left(\hat{g}_m(x, \pi_{\hat{g}_m}(x)) - \hat{g}_m(x, a) \right)} \right] \leq \frac{K}{\gamma_m}. \end{aligned}$$

□

Lemma 4 states that the expected inverse probability is bounded by the estimated expected instantaneous regret of policy π .

Lemma 4. For any epoch m and any policy $\pi \in \Psi$,

$$V(p_m, \pi) \leq K + \gamma_m \text{Reg}_{\hat{g}_m}(\pi).$$

Proof. For any policy $\pi \in \Psi$, given any context $x \in \mathcal{X}$,

$$\frac{1}{p_m(\pi(x)|x)} \begin{cases} = K + \gamma_m (\hat{g}_m(x, \hat{a}_m(x)) - \hat{g}_m(x, \pi(x))) & \text{if } \pi(x) \neq \hat{a}_m(x) \\ \leq \frac{1}{1/K} = K = K + \gamma_m (\hat{g}_m(x, \hat{a}_m(x)) - \hat{g}_m(x, \pi(x))) & \text{if } \pi(x) = \hat{a}_m(x). \end{cases}$$

Thus,

$$V(p_m, \pi) = \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\frac{1}{p_m(\pi(x)|x)} \right] \leq K + \gamma_m \mathbb{E}_{x \sim D_{\mathcal{X}}} [\hat{g}_m(x, \hat{a}_m(x)) - \hat{g}_m(x, \pi(x))] = K + \gamma_m \text{Reg}_{\hat{g}_m}(\pi).$$

□

B.7 Bounding Treatment Effect Error

For any policy, Lemma 5 bounds the error of the treatment effect estimate for the first few epochs. Our definition of m^* and our choice of γ_{m+1} allows us to prove this lemma without assuming realizability.

Lemma 5. Suppose the event \mathcal{W}_1 defined in (26) holds. Then, for any policy $\pi \in \Psi$ and epoch $m \leq \min\{m^*, \hat{m}\} + 1$,

$$|\Phi_{p_m}^{\hat{g}_m}(\pi) - \Phi_{p_m}(\pi)| \leq \frac{\sqrt{V(p_m, \pi)} \sqrt{K}}{2\gamma_{m+1}}.$$

Proof. For any policy π and epoch $m \leq \min\{m^*, \hat{m}\} + 1$,

$$\begin{aligned}
& |\Phi_{p_m}^{\hat{g}_{m+1}}(\pi) - \Phi_{p_m}(\pi)| \\
& \leq \mathbb{E}_{x \sim D_{\mathcal{X}}} [|\langle e_{\pi(x)} - p_m(x), \hat{g}_{m+1}(x) - g^*(x) \rangle|] \\
& = \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\sqrt{\frac{1}{p_m(\pi(x)|x)} p_m(\pi(x)|x) \langle e_{\pi(x)} - p_m(x), \hat{g}_{m+1}(x) - g^*(x) \rangle^2} \right] \\
& \leq \mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\sqrt{\frac{1}{p_m(\pi(x)|x)} \mathbb{E}_{a \sim p_m(\cdot|x)} [\langle e_a - p_m(x), \hat{g}_{m+1}(x) - g^*(x) \rangle^2]} \right] \\
& \leq \sqrt{\mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\frac{1}{p_m(\pi(x)|x)} \right]} \sqrt{\mathbb{E}_{x \sim D_{\mathcal{X}}} \mathbb{E}_{a \sim p_m(\cdot|x)} [\langle e_a - p_m(x), \hat{g}_{m+1}(x) - g^*(x) \rangle^2]} \\
& = \sqrt{\mathbb{E}_{x \sim D_{\mathcal{X}}} \left[\frac{1}{p_m(\pi(x)|x)} \right]} \sqrt{\mathcal{E}_{p_m}(\hat{g}_{m+1})} \\
& \leq \sqrt{V(p_m, \pi)} \sqrt{2\xi \left(\tau_m - \tau_{m-1}, \frac{\delta'}{m^2} \right)} \\
& = \frac{\sqrt{V(p_m, \pi)} \sqrt{K}}{2\gamma_{m+1}}.
\end{aligned}$$

The first inequality follows from Jensen's inequality, the second inequality is straight-forward, the third inequality follows from Cauchy-Schwarz inequality, and the last inequality follows from assuming that \mathcal{W}_1 from (26) holds. \square

B.8 Bounding Implicit Regret Errors and True Regret

The next lemma implies that before misspecification becomes a problem we are able to bound regret. We omit the proof since it is very similar to the proof of Lemma 8 in Simchi-Levi and Xu (2021), given how we defined the instantaneous treatment effect and instantaneous regret. The proof of this result relies on an inductive argument over the epochs and making use of Lemmas 4 and 5.

Lemma 6. *Suppose the event \mathcal{W}_1 defined in (26) holds. Let $C_1 = 5.15$. For any policy $\pi \in \Psi$ and epoch $m \leq \min\{m^*, \hat{m}\} + 1$,*

$$\begin{aligned}
\text{Reg}(\pi) & \leq 2\text{Reg}_{\hat{g}_m}(\pi) + \frac{C_1 K}{\gamma_m} \\
\text{Reg}_{\hat{g}_m}(\pi) & \leq 2\text{Reg}(\pi) + \frac{C_1 K}{\gamma_m}
\end{aligned}$$

In other words, for any policy, Lemma 6 bounds the error of the regret estimate for the first few epochs. Lemma 7 bounds the implicit regret of the randomized policy Q_m for the first few epochs. Lemma 7 and its proof is more or less the same as the statement and the proof of Lemma 9 in Simchi-Levi and Xu (2021).

Lemma 7. *Suppose the event \mathcal{W}_1 defined in (26) holds. Then, for any epoch $m \leq \min\{m^*, \hat{m}\} + 1$,*

$$\sum_{\pi \in \Psi} Q_m(\pi) \text{Reg}(\pi) \leq \frac{(2 + C_1)K}{\gamma_m}.$$

Proof. For any $m \leq \min\{m^*, \hat{m}\} + 1$:

$$\sum_{\pi \in \Psi} Q_m(\pi) \text{Reg}(\pi) \leq \sum_{\pi \in \Psi} Q_m(\pi) \left(2\text{Reg}_{\hat{g}_m}(\pi) + \frac{C_1 K}{\gamma_m} \right) \leq \frac{2K}{\gamma_m} + \frac{C_1 K}{\gamma_m},$$

where the first inequality follows from Lemma 6, and the second inequality follows from Lemma 3. \square

B.9 Proof of Theorem 1

Proof. For each round $t \leq T$, define

$$M_t := r_t(\pi^*(x)) - r_t(a_t) - \sum_{\pi \in \Psi} Q_{m(t)}(\pi) \text{Reg}(\pi).$$

By Lemma 2, we have $\mathbb{E}[M_t | \Gamma_{t-1}] = 0$. Since $|M_t| \leq 2$, M_t is a martingale difference sequence. By Azuma's inequality,

$$\sum_{t=1}^T M_t \leq \sqrt{8T \log(2/\delta)}$$

with probability at least $1 - \delta/2$. In other words,

$$\sum_{t=1}^T (r_t(\pi^*(x_t)) - r_t(a_t)) \leq \sum_{t=1}^T \sum_{\pi \in \Psi} Q_{m(t)}(\pi) \text{Reg}(\pi) + \sqrt{8T \log(2/\delta)}$$

with probability at least $1 - \delta/2$. We consider two cases, when $T \leq \tau_{m^*+1}$ and $T > \tau_{m^*+1}$.

Case 1 ($T \leq \tau_{m^*+1}$): If $T \leq \tau_{m^*+1}$, then the status of the algorithm is always safe by Assumption 2. Moreover, by Lemma 1, with probability at least $1 - \delta/2$, the event \mathcal{W}_1 holds. Therefore, by a union bound and then using Lemma 7, we bound

$$\begin{aligned} \sum_{t=1}^T (r_t(\pi^*(x_t)) - r_t(a_t)) &\leq \sqrt{8T \log(2/\delta)} + \sum_{t=1}^T \sum_{\pi \in \Psi} Q_{m(t)}(\pi) \text{Reg}(\pi) \\ &\leq \sqrt{8T \log(2/\delta)} + \tau_1 + \sum_{t=\tau_1+1}^T \frac{(2+C_1)K}{\gamma_{m(t)}} \\ &\leq \sqrt{8T \log(2/\delta)} + \tau_1 + (2+C_1) \sum_{t=\tau_1+1}^T \sqrt{8K \cdot \xi \left(\tau_{m(t)-1} - \tau_{m(t-1)}, \frac{\delta}{2m(t)^2} \right)} \\ &\leq \mathcal{O} \left(\sqrt{K} \sum_{t=\tau_1+1}^T \sqrt{8\xi \left(\tau_{m(t)-1} - \tau_{m(t-1)}, \frac{\delta}{2m(t)^2} \right)} \right) \end{aligned}$$

with probability at least $1 - \delta$.

Case 2 ($T > \tau_{m^*+1}$): For the cumulative regret up to τ_{m^*+1} , we can follow the arguments of case 1. For the cumulative regret between $\tau_{m^*+1} + 1$ and T , we use Assumption 3 to bound expected regret in terms of the misspecification error at these rounds. By a union bound, we have

$$\begin{aligned} \sum_{t=1}^T (r_t(\pi^*(x_t)) - r_t(a_t)) &\leq \sqrt{8T \log(2/\delta)} + \sum_{t=1}^T \sum_{\pi \in \Psi} Q_{m(t)}(\pi) \text{Reg}(\pi) \\ &= \sqrt{8T \log(2/\delta)} + \sum_{t=1}^{\tau_{m^*+1}} \sum_{\pi \in \Psi} Q_{m(t)}(\pi) \text{Reg}(\pi) + \sum_{t=\tau_{m^*+1}+1}^T \sum_{\pi \in \Psi} Q_{m(t)}(\pi) \text{Reg}(\pi) \\ &\leq \sqrt{8T \log(2/\delta)} + \tau_1 + (2+C_1) \sum_{t=\tau_1+1}^{\tau_{m^*+1}} \sqrt{8K \cdot \xi \left(\tau_{m(t)-1} - \tau_{m(t-1)}, \frac{\delta}{2m(t)^2} \right)} \\ &\quad + (T - \tau_{m^*+1} - 1) \mathcal{O}(\sqrt{KB}) \\ &\leq \mathcal{O} \left(\sqrt{K} \sum_{t=\tau_1+1}^T \sqrt{\xi \left(\tau_{m(t)-1} - \tau_{m(t-1)}, \frac{\delta}{2m(t)^2} \right)} + \sqrt{KBT} \right). \end{aligned}$$

with probability at least $1 - \delta$. □

C UNKNOWN PROPENSITIES

For completeness, consider the case where we also do not know the exact treatment propensities. It becomes necessary to estimate the treatment propensity nuisance parameter in cases where the treatment assignment probabilities are not logged by the algorithm or when there is treatment non-compliance from the users. In these settings, the results may depend on the estimation of the nuisance parameters.

Let \hat{p} and $\hat{\mu}$ be the estimates of p and μ , respectively. Consider the approximate R -loss risk under kernel p with the estimated nuisance propensities \hat{p} in the loss function

$$\mathcal{R}_{p,\hat{p}}(g) := \mathbb{E}_{(x,a,r(a)) \sim D(p)} \left[\left(r(a) - \hat{\mu}(x) - \langle e_a - \hat{p}(x), g(x) \rangle \right)^2 \right]$$

and its corresponding excess risk

$$\mathcal{E}_{p,\hat{p}}(g) := \mathcal{R}_{p,\hat{p}}(g) - \min_{g'} \mathcal{R}_{p,\hat{p}}(g'),$$

where the infimum is over all functions. Note how $\mathcal{R}_{p,p}(g) = \mathcal{R}_p(g)$ and $\mathcal{E}_{p,p}(g) = \mathcal{E}_p(g)$ from our previous definitions. We aim to establish *true* R -loss excess risk bounds in the setting where we also have to estimate the treatment propensities. First, observe that the risk under the estimated kernel \hat{p} is

$$\begin{aligned} \mathcal{R}_{p,\hat{p}}(g) &= \mathbb{E} \left[\left(r(a) - \hat{\mu}(x) - \langle e_a - \hat{p}(x), g(x) \rangle \right)^2 \right] \\ &= \mathbb{E} \left[\left(r(a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle - \langle p(x) - \hat{p}(x), g(x) \rangle \right)^2 \right] \\ &= \mathbb{E} \left[\left(r(a) - \hat{\mu}(x) - \langle e_a - p(x), g(x) \rangle \right)^2 \right] \\ &\quad + \mathbb{E} \left[\langle p(x) - \hat{p}(x), g(x) \rangle^2 - 2(\mu(x) - \hat{\mu}(x)) \langle p(x) - \hat{p}(x), g(x) \rangle \right] \\ &= \mathcal{R}_p(g) + \mathbb{E} \left[\langle p(x) - \hat{p}(x), g(x) \rangle^2 - 2(\mu(x) - \hat{\mu}(x)) \langle p(x) - \hat{p}(x), g(x) \rangle \right]. \end{aligned}$$

Recall $\mathcal{G}^* := \arg \min_{g'} \mathcal{R}_p(g')$ is the set of true treatment effects models under the R -loss risk with the true kernel p . For any $g^* \in \mathcal{G}^*$,

$$\mathcal{R}_{p,\hat{p}}(g) - \mathcal{R}_{p,\hat{p}}(g^*) = \mathcal{E}_p(g) + \mathbb{E} \left[\left(\langle p(x) - \hat{p}(x), g(x) + g^*(x) \rangle - 2(\mu(x) - \hat{\mu}(x)) \right) \cdot \langle p(x) - \hat{p}(x), g(x) - g^*(x) \rangle \right]$$

and, by Cauchy-Schwarz inequality,

$$\begin{aligned} |\mathcal{R}_{p,\hat{p}}(g) - \mathcal{R}_{p,\hat{p}}(g^*) - \mathcal{E}_p(g)| &= \left| \mathbb{E} \left[\left(\langle p(x) - \hat{p}(x), g(x) + g^*(x) \rangle - 2(\mu(x) - \hat{\mu}(x)) \right) \cdot \langle p(x) - \hat{p}(x), g(x) - g^*(x) \rangle \right] \right| \\ &\leq \mathbb{E} \left[\left(\|p(x) - \hat{p}(x)\| \|g(x) + g^*(x)\| + 2|\mu(x) - \hat{\mu}(x)| \right) \cdot \|p(x) - \hat{p}(x)\| \|g(x) - g^*(x)\| \right] \\ &\leq 2\sqrt{K} \mathbb{E} \left[\|g(x) - g^*(x)\|^2 \right]^{1/2} \cdot \mathbb{E} \left[\|p(x) - \hat{p}(x)\|^4 \right] \\ &\quad + 2 \mathbb{E} \left[\|g(x) - g^*(x)\|^2 \right]^{1/2} \cdot \mathbb{E} \left[|\mu(x) - \hat{\mu}(x)|^2 \|p(x) - \hat{p}(x)\|^2 \right]^{1/2} \\ &= C_2 \mathbb{E} \left[\|g(x) - g^*(x)\|^2 \right]^{1/2} \cdot \left(\mathbb{E} \left[\|p(x) - \hat{p}(x)\|^4 \right]^{1/2} \right. \\ &\quad \left. + \mathbb{E} \left[|\mu(x) - \hat{\mu}(x)|^2 \|p(x) - \hat{p}(x)\|^2 \right]^{1/2} \right) \\ &:= \Delta(g, g^*), \end{aligned}$$

where $C_2 = 2(1 + \sqrt{K}) \leq 4\sqrt{K}$. Since $\mathcal{R}_{p,\hat{p}}(g^*) = \mathcal{R}_{p,\hat{p}}(f^*)$, the quantity on the left-hand side of the inequality above does not depend on the choice of $g^* \in \mathcal{G}^*$ while the bound on the right-hand side does. Therefore, we can take the minimum over \mathcal{G}^* on the right-hand side of the above inequality to get

$$|\mathcal{R}_{p,\hat{p}}(g) - \mathcal{R}_{p,\hat{p}}(g^*) - \mathcal{E}_p(g)| \leq \min_{g^* \in \mathcal{G}^*} \Delta(g, g^*).$$

Now, notice that

$$\mathcal{E}_p(g) \leq \mathcal{R}_{p,\hat{p}}(g) - \mathcal{R}_{p,\hat{p}}(g^*) + |\mathcal{R}_{p,\hat{p}}(g) - \mathcal{R}_{p,\hat{p}}(g^*) - \mathcal{E}_p(g)| \leq \mathcal{E}_{p,\hat{p}}(g) + \min_{g^* \in \mathcal{G}^*} \Delta(g, g^*). \quad (29)$$

The first term $\mathcal{E}_{p,\hat{p}}(g)$ on the right-hand side of the last inequality is the approximate excess risk with the estimated kernel \hat{p} . This can be bounded by the excess risk estimation rates for the R -oracle that uses the estimated kernel \hat{p} . Therefore, to bound the excess risk under the true kernel, we need to bound the second term $\min_{g^* \in \mathcal{G}^*} \Delta(g, g^*)$. This requires us to bound the following:

$$\begin{aligned} \min_{g^* \in \mathcal{G}^*} \mathbb{E}[\|g(x) - g^*(x)\|^2] &\leq \min_{g^* \in \mathcal{G}^*} \frac{1}{\min_{x,a} p(a|x)} \mathbb{E} \left[\sum_a p(a|x) (g(x,a) - g^*(x,a))^2 \right] \\ &\leq \frac{1}{\min_{x,a} p(a|x)} \min_{g^* \in \mathcal{G}^*} \left(\mathcal{E}_p(g) + \mathbb{E}_{x \sim D_X} \left[\mathbb{E}_{a \sim p(\cdot|x)} [g(x,a) - g^*(x,a)]^2 \right] \right), \end{aligned}$$

where the last inequality follows from the result derived in Appendix A.2. Furthermore, note that for a fixed $g \in \mathcal{G}$, there exists a $g^* \in \mathcal{G}^*$ such that $\mathbb{E}_{a \sim p(\cdot|x)} [g^*(x,a)] = \mathbb{E}_{a \sim p(\cdot|x)} [g(x,a)]$. In particular, since $f^* \in \mathcal{G}^*$, the function $g^*(x,a) = f^*(x,a) + \mathbb{E}_{a \sim p(\cdot|x)} [f^*(x,a) - g(x,a)]$ is clearly in \mathcal{G}^* and it satisfies this condition. Therefore, the second term in the above inequality vanishes when the infimum over \mathcal{G}^* is taken and we get that

$$\min_{g^* \in \mathcal{G}^*} \mathbb{E}[\|g(x) - g^*(x)\|^2] \leq \frac{\mathcal{E}_p(g)}{\min_{x,a} p(a|x)} \leq \frac{\mathcal{E}_p(g)}{\eta} \quad (30)$$

where $\eta := \min_{x,a} p(a|x)$ is the minimum overlap of kernel p . Now, suppose we can estimate the nuisance parameters with n data samples such that

$$\begin{aligned} \mathbb{E}[\|p(x) - \hat{p}(x)\|^4] &\leq K\theta_n \\ \mathbb{E}[|\mu(x) - \hat{\mu}(x)|^2 \|p(x) - \hat{p}(x)\|^2] &\leq K\theta_n \end{aligned}$$

for some rate $\theta_n = \mathcal{O}(n^{-\kappa})$ with $\kappa \geq \frac{1}{2}$. Note that

$$\begin{aligned} \mathbb{E}[\|p(x) - \hat{p}(x)\|^4] &\leq K \mathbb{E}[\|p(x) - \hat{p}(x)\|^2] \\ \mathbb{E}[|\mu(x) - \hat{\mu}(x)|^2 \|p(x) - \hat{p}(x)\|^2] &\leq K \mathbb{E}[|\mu(x) - \hat{\mu}(x)|^2]. \end{aligned}$$

Therefore, we can estimate the nuisance parameters using standard least squares estimation procedures to guarantee fast estimation rates, i.e., $\kappa = 1$, for the above quantities. By combining result (30) and the above estimation rate assumptions on the nuisance parameters into inequality (29), we get

$$\mathcal{E}_p(\hat{g}) \leq \mathcal{E}_{p,\hat{p}}(\hat{g}) + \sqrt{\frac{C_3 K \theta_n \mathcal{E}_p(\hat{g})}{\eta}},$$

where $C_3 = 64$. To form bounds on $\mathcal{E}_p(\hat{g})$, we consider the following two cases.

Case 1:

$$\sqrt{\frac{C_3 K \theta_n}{\eta}} \leq \frac{\sqrt{\mathcal{E}_p(\hat{g})}}{2}$$

In this case, we have

$$\mathcal{E}_p(\hat{g}) \leq 2\mathcal{E}_{p,\hat{p}}(\hat{g}) \leq \mathcal{O}(\rho_n),$$

where we let ρ_n be the rate on the approximate R -loss excess risk with the estimated kernel. Similar results discussed in Appendix A.3 can be used to establish fast rates for this excess risk. Therefore, the excess risk rate in this domain is controlled by the fast rates on the approximate R -loss excess risk with the estimated kernel. This is independent of the quality of nuisance parameter estimation. In this regime, the bandit algorithm can have an aggressive exploration rate.

Case 2:

$$\frac{\sqrt{\mathcal{E}_p(\hat{g})}}{2} < \sqrt{\frac{C_3 K \theta_n}{\eta}}$$

In this case,

$$\mathcal{E}_p(\hat{g}) < \frac{4C_3 K \theta_n}{\eta}.$$

In this domain, the excess risk rate is controlled by the orthogonal nuisance parameter estimation rates. Moreover, in this setting, overlap comes into play. In our bandit setting, the overlap parameter η is vanishing as well. The rate of decrease of kernel overlap is dictated by the excess risk estimation rate. In particular, using the excess risk rate in this case gives $\eta = \mathcal{O}(1/\gamma_n) = \mathcal{O}(\sqrt{\theta_n/\eta})$. This requires $\eta = \mathcal{O}(\theta_n^{1/3})$ which implies

$$\mathcal{E}_p(\hat{g}) \leq \mathcal{O}(K\theta_n^{2/3}).$$

Therefore, we find that overlap helps with orthogonality. However, it diminishes the required rate of exploration in the algorithm since the nuisance parameters also have fast rate guarantees. In this regime, the bandit algorithm is not as aggressive in its exploration. If we wanted to adapt the algorithm to these different regimes, we would need to characterize when the different estimation rates dominate kick in.

D EXPERIMENTAL SETUP

We used the EconML package (Battocchi et al., 2019) for the R -learner implementation on a linear model class and we used the VowpalWabbit Cobra package (Coba, 2022) for running our synthetic contextual bandit benchmark simulations. Refer to the Github repositories for package licenses. A single 2.8 GHz Quad-Core Intel Core i7 CPU was used to run the synthetic experiments. The code is available at <https://github.com/agcarranza/hte-bandits>.