
Causal Q-Aggregation for CATE Model Selection

Hui Lan

Stanford University
huilan@stanford.edu

Vasilis Syrgkanis

Stanford University
vsyrgk@stanford.edu

Abstract

Accurate estimation of conditional average treatment effects (CATE) is at the core of personalized decision making. While there is a plethora of models for CATE estimation, model selection is a nontrivial task, due to the fundamental problem of causal inference. Recent empirical work provides evidence in favor of proxy loss metrics with double robust properties and in favor of model ensembling. However, theoretical understanding is lacking. Direct application of prior theoretical work leads to suboptimal oracle model selection rates due to the non-convexity of the model selection problem. We provide regret rates for the major existing CATE ensembling approaches and propose a new CATE model ensembling approach based on Q-aggregation using the doubly robust loss. Our main result shows that causal Q-aggregation achieves statistically optimal oracle model selection regret rates of $\frac{\log(M)}{n}$ (with M models and n samples), with the addition of higher-order estimation error terms related to products of errors in the nuisance functions. Crucially, our regret rate does not require that any of the candidate CATE models be close to the truth. We validate our new method on many semi-synthetic datasets and also provide extensions of our work to CATE model selection with instrumental variables and unobserved confounding.

1 INTRODUCTION

Identifying optimal decisions requires understanding the causal effect of an action on an outcome of inter-

est. With the emergence of rich and large datasets in many application domains such as digital experimentation, precision medicine, and digital marketing, identifying optimal personalized decisions has emerged as a mainstream topic in the literature. Identifying optimal personalized decisions requires understanding how the causal effect changes with observable characteristics of the treated unit. For this reason, many recent works have studied the estimation of conditional average treatment effects (CATE):

$$\tau_0(X) := \mathbb{E}[Y(1) - Y(0) \mid X]$$

where X are observable features and $Y(d)$ is the potential outcome of interest under treatment $d \in \{0, 1\}$.

This has led to a surge of many different methods for CATE estimation using machine learning techniques, such as deep learning (Shalit et al., 2017; Shi et al., 2019), lasso (Imai and Ratkovic, 2013), random forests (Wager and Athey, 2018; Oprescu et al., 2019), Bayesian regression trees (Hahn et al., 2020), as well as model-agnostic frameworks such as meta-learners (Künzel et al., 2019) and double machine learning (Kennedy, 2020; Foster and Syrgkanis, 2023; Nie and Wager, 2021). Machine learning CATE estimation has been considered both under the assumption of unconfoundedness, i.e., that outcomes $Y(1), Y(0)$ are independent of the assigned treatment D , conditional on the observed features X , as well as when there is unobserved confounding and access to instrumental variables (Hartford et al., 2017; Syrgkanis et al., 2019).

However, the effectiveness of each approach depends on the learnability of the various causal mechanisms at play and the structure of the CATE function. This has led to the pursuit of automated and data-driven model selection approaches for CATE estimation (Schuler et al., 2018; Mahajan et al., 2022; Curth and van der Schaar, 2023; Alaa and Van Der Schaar, 2019). Due to the fundamental problem of causal inference, i.e., that we do not observe both counterfactual outcomes $Y(1), Y(0)$ for each unit, there does not exist a perfect analogue of out-of-sample model selection based on mean squared error, as is typically invoked in regression problems. For this reason, many works have

proposed proxy loss metrics Nie and Wager (2021); Foster and Syrgkanis (2023); Kennedy (2020); Alaa and Van Der Schaar (2019), analogous to the mean square error, which, under assumptions, can be used for data-driven model evaluation. In particular, Mahajan et al. (2022), highlights that proxy metrics that incorporate double robustness properties (that we either have a good estimate of the outcome regression function $\mathbb{E}[Y | D, X]$ or the propensity function $\mathbb{E}[D | X]$), tend to perform well in many practical scenarios.

Moreover, ensembling of CATE models tends to typically increase performance, as compared to simply choosing the single best model based on the loss. The empirical success of model aggregation in causal inference has also been demonstrated multiple times. In the 2016 Atlantic Causal Inference Conference Competition, 3 out of the 5 best performing models to estimate the average treatment effect were ensemble models (Dorie et al., 2019). Moreover, Athey et al. (2019) also showed that ensemble methods can outperform single models at predicting counterfactuals.

Contributions and prior work. Motivated by these empirical findings and by the success of ensemble methods in the regression and supervised machine learning literature, we consider the problem of CATE model selection: given a set of M candidate CATE models $\mathcal{F} = \{f_1, \dots, f_M\}$ and n samples, can we construct an ensemble $\hat{f} = \sum_j \theta_j f_j$ that competes with the best model:

$$\mathbb{E}[(f_{\hat{\theta}}(X) - \tau_0(X))^2] \leq \min_{j=1}^M \mathbb{E}[(f_j(X) - \tau_0(X))^2] + \epsilon$$

No rigorous theoretical treatment of this CATE model selection problem from observational data has been provided in the literature. Prior theoretical work considered only randomized trial data Han and Wu (2022), or considered only CATE estimation over rich function spaces Nie and Wager (2021); Foster and Syrgkanis (2023), which is technically very different from the goal of competing with the best of M given models, or consider only guarantees for model evaluation Alaa and Van Der Schaar (2019), but not selection. For the case of regression problems, the optimal oracle selection error ϵ that is achievable is of the order of $\log(M)/n$. This rate allows selecting from an exponentially sized set of candidate models M and decays fast with the sample size, as n^{-1} instead of $n^{-1/2}$.

Our main result is such an optimal oracle selection error for CATE model selection. We propose a novel CATE model ensembling method based on Q-aggregation Lecué and Rigollet (2014) with a doubly robust loss and prove that it attains the oracle selection error of $\log(M)/n$ plus second-order terms that

depend on the regression and the propensity estimation errors. Moreover, even the $\log(M)$ term can be removed if one has a relatively good prior over which models will be the better performers, which can be quite natural in the context of CATE model selection.

In the regression setting, optimal selection error rates cannot be achieved using ERM approaches, i.e., by minimizing squared loss over some ensemble hypothesis space Lecué and Rigollet (2014). Thus, prior approaches to CATE ensembles Nie and Wager (2021); Han and Wu (2022) that either select the best model based on a proxy loss, or optimize over the space of all possible convex combinations of models, will necessarily achieve suboptimal selection error rates (either $\sqrt{\log(M)/n}$ or M/n in the worst case). As a set of side results, we provide oracle selection rates for both of these methods for the CATE problem (which are novel in the literature and require a generalization of the main statistical learning theorems with nuisance functions presented in Foster and Syrgkanis (2023)). These results also highlight when one should expect ERM approaches to achieve close to optimal error.

Moreover, a theoretical result on optimal oracle model selection necessarily needs to use techniques beyond what is outlined in the prior theoretical work of Nie and Wager (2021); Foster and Syrgkanis (2023); Kennedy (2020), which is our main technical contribution. Notably, the general theorems of Foster and Syrgkanis (2023) require a first-order optimality condition, which is not satisfied due to the non-convexity of the function space \mathcal{F} with which we are trying to compete and cannot be applied to prove our main Q-aggregation theorem. Due to space constraints, we defer a detailed comparison to previous work in the Appendix A.

Finally, we provide an extensive empirical evaluation of our causal Q-aggregation ensemble method and of the ERM based ensemble approaches, on a broad spectrum of synthetic data as well as semi-synthetic data based on five real-world datasets from a multitude of application domains, such as political science (Green and Kern, 2012), economics (Farbmacher et al., 2021; Poterba and Venti, 1994; Poterba et al., 1995), education Word et al. (1990) and digital advertising (Diemert Eustache, Betlei Artem et al., 2018). We showcase how ensemble methods offer more robust solutions than any single meta-learning approach proposed in the literature and how Q-aggregation can many times lead to improvements in performance than ERM based approaches.

2 PRELIMINARIES

We consider a setting where we are given a set of M candidate CATE models $\{f_1, \dots, f_M\}$ and a data set of n samples. Each candidate model f_i maps covariates X to a CATE prediction and can be thought of as a CATE model that was fitted on a separate set of samples (e.g. fitted on training set and correspond to different hyperparameters, methods, and model spaces), and the n samples that we are using are the validation set. Our goal is to select a model or an ensemble of models, i.e. $f_\theta = \sum_{j=1}^M \theta_j f_j$, that achieves mean-squared-error with respect to the true CATE function $\tau_0(x) = \mathbb{E}[Y(1) - Y(0) \mid X = x]$ that is comparable to the best model, i.e.:

$$\mathbb{E}[(f_\theta(X) - \tau_0(X))^2] \leq \min_{j=1}^M \mathbb{E}[(f_j(X) - \tau_0(X))^2] + \epsilon$$

As we will see in Section 5, the true CATE model τ_0 , can be expressed as the minimizer of a risk $\min_f \mathbb{E}[\ell(Z; f(X), g_0)]$, which, in addition to the CATE function f also depends on other nuisance functions g_0 , which are the solution to auxiliary statistical estimation problems and which also need to be estimated from data. Moreover, the loss function varies depending on the argument that is used to identify the CATE (e.g. conditional exogeneity, instrumental variables, etc.).

For this reason, we will consider the more general problem of model selection in the presence of nuisance functions and present our main theorems at this level of generality. Subsequently in Section 5 and Appendix 6, we will instantiate the general theorems for particular CATE selection problems.

Model Selection with Nuisance Functions

Given a set of candidate functions $\mathcal{F} = \{f_1, \dots, f_M\}$, we care about minimizing an expected loss function $\mathbb{E}[\ell(Z; f(X), g_0)]$, where g_0 is the ground truth nuisance function. In particular, our goal is to identify an ensemble of the candidate estimators, i.e.

$$f_\theta := \sum_{j=1}^M \theta_j f_j, \quad \theta \in \Theta := \left\{ \theta \in \mathbb{R}_{\geq 0}^M : \sum_{j=1}^M \theta_j = 1 \right\}$$

that controls the oracle selection error:

$$\mathcal{S}(\theta) := \mathbb{E}[\ell(Z; f_\theta(X), g_0)] - \min_{j=1}^M \mathbb{E}[\ell(Z; f_j(X), g_0)]$$

Our target is to prove an oracle inequality of the order: $O\left(\frac{\log(M)}{n}\right)$, with access to n samples of the variables (X, Z) , which is the optimal rate in the absence of nuisance functions. Since, we also do not know g_0 , we

need to construct an estimate \hat{g} from the data. Our goal is to derive an error that would have a second order dependence on the estimation error of g_0

$$\mathcal{S}(\theta) \leq O\left(\frac{\log(M)}{n}\right) + \text{ERROR}(\hat{g})$$

for some appropriately defined error function that would decay at least as fast or faster than $1/n$ under assumptions on the estimation error of \hat{g} . *Unless otherwise stated, the estimate \hat{g} will be assumed to be trained on a separate sample of size n (e.g. the training sample used to fit the models f_i).*

Notation We use short-hand notation:

$$\begin{aligned} P\ell(Z; f(X), g) &= \mathbb{E}[\ell(Z; f(X), g)] \\ P_n\ell(Z; f(X), g) &= \frac{1}{n} \sum_i \ell(Z_i; f(X_i), g), \end{aligned}$$

where Z denotes all the random variables and X is some subset of them. Moreover, we denote with $\|f - f'\|^2 = P(f'(X) - f(X))^2$ and with $\|X\|_{L^p} = \mathbb{E}[X^p]^{1/p}$. For any convex combination θ over a set of candidate functions $\{f_1, \dots, f_M\}$, we define the notation:

$$\ell_{\theta, g}(Z) := \ell(Z; f_\theta(X), g) \quad R(\theta, g) := P\ell_{\theta, g}(Z)$$

We also denote with $e_j \in \Theta$ the vector that has 1 at coordinate j and zero otherwise (that is, $f_{e_j} = f_j$) and denote with $\text{co}(\mathcal{F})$ the convex hull of space \mathcal{F} .

In addition, We let $\hat{Y}(g)$ denote the debiased pseudo-outcome:

$$\hat{Y}(g) = h(1, W) - h(0, W) + a(D, W)(Y - h(D, W))$$

with $g = (h, a)$, h an estimate of $h_0(D, X) = \mathbb{E}[Y \mid D, X]$ and a an estimate of the signed inverse propensity function a_0 (a.k.a. the Riesz representer Chernozhukov et al. (2022)):

$$a_0(D, W) := \frac{D - p_0(W)}{p_0(W)(1 - p_0(W))}, \quad p_0(W) := \mathbb{E}[D \mid W]$$

3 ERM-BASED SELECTION

We start our investigation by exploring what regret rates are achievable by ensembling or stacking approaches based on ERM, i.e. minimizing the empirical risk over the weight parameters $\theta \in \Theta$, for some constraint space Θ . To simplify exposition, we restrict attention in this section to square losses:

$$\ell(Z; f(X), g) = (\hat{Y}(g) - f(X))^2 \quad (1)$$

where g is a vector of nuisance functions and $\hat{Y}(g)$ is some functional of these functions that satisfies:

$$\tau_0(X) := \mathbb{E}[\hat{Y}(g_0) \mid X] \quad (2)$$

for some target function τ_0 .

Two major ensemble methods are choosing the single best model, $(\Theta_b = \{e_1, \dots, e_M\})$, and choosing the best convex combination,

$$\hat{\theta}_b := \arg \min_{\theta \in \Theta_b} \mathbb{E}_n[(\hat{Y}(\hat{g}) - f_\theta)^2] \quad (\text{Best-ERM})$$

$$\hat{\theta}_c := \arg \min_{\theta \in \Theta} \mathbb{E}_n[(\hat{Y}(\hat{g}) - f_\theta)^2] \quad (\text{Convex-ERM})$$

It is known that such ERM approaches lead to suboptimal regret rates without further assumptions, even in the absence of nuisance functions (i.e., when we know g_0) Lecué and Rigollet (2014). As a side result of this paper, we provide oracle results for these estimators in the presence of nuisance functions and refine prior results in the regression literature, providing some insights as to when we should expect these methods to achieve near optimal performance in practice.

Theorem 1 (Guarantees for Best-ERM). *Consider the case of the square loss and let:*

$$\text{BIAS}(X; \hat{g}) = \mathbb{E} \left[\hat{Y}(\hat{g}) - \hat{Y}(g_0) \mid X \right]$$

Assume that $\hat{Y}(g)$ is absolutely uniformly bounded by U for all $g \in G$ and all functions f_j are absolutely uniformly bounded by U . Let

$$\bar{f}_* := \arg \min_{f \in \text{co}(\mathcal{F})} \|f - \tau_0\|, \quad f_* := \arg \min_{f \in \mathcal{F}} \|f - \tau_0\|$$

and let:

$$\epsilon = \|f_* - \tau_0\|^2 - \|\bar{f}_* - \tau_0\|^2$$

Then with probability $1 - \delta$:

$$\mathcal{S}(\hat{\theta}_b) \leq O \left(\epsilon + \frac{\log((M \vee n)/\delta)}{n} + \mathbb{E}[\text{BIAS}(X; \hat{g})^2] \right)$$

The result is enabled by a generalization of the main theorems in Foster and Syrgkanis (2023). Incorporating the bound $\|f_* - \tau_0\|^2 - \|\bar{f}_* - \tau_0\|^2$ is a novel addition to the literature, even in the absence of nuisance functions. We see that the best selector achieves the target rate only when ϵ vanishes faster than $\log(M)/n$, which would happen only if the best model in \mathcal{F} is very close to the true model τ_0 (c.f. $\min_{f \in \mathcal{F}} \|f - \tau_0\|^2 \leq \epsilon$), which is quite unlikely when τ_0 is complex, or when the best model \bar{f}_* in the convex hull ends up putting most of the support on only one model, i.e., there is a dominant model (c.f. $\min_{f \in \mathcal{F}} \|f - f_*\| \leq \epsilon$). Typical prior bounds depend on the larger quantity $\epsilon = \|f_* - \tau_0\|^2$, which is not close to zero if the optimizer in the convex hull is primarily supported on one model.

Theorem 1 is also closely related to the results in Van Der Laan and Dudoit (2003). As we show in Appendix A.1, our Theorem 1 improves upon this prior

work in: i) only incurring a dependence on the nuisance error through the term $\mathbb{E}[\text{BIAS}(X; \hat{g})^2]$, which possesses doubly robust properties in our main application of CATE model selection, unlike the bound in Van Der Laan and Dudoit (2003), which would also depend on a term of the form $\mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2]$, which does not possess doubly robust properties, ii) only depending on the smaller quantity ϵ , unlike the bound in Van Der Laan and Dudoit (2003) which would depend on $\|f_* - \tau_0\|^2$; as we discussed ϵ can be much smaller than $\|f_* - \tau_0\|^2$, when there is a dominant model, but all models are far from τ_0 .

Theorem 2 (Guarantees for Convex-ERM). *Consider the case of the square loss and assume that $\hat{Y}(g)$ is absolutely uniformly bounded by U for all $g \in G$ and all functions f_j are absolutely uniformly bounded by U . Then with probability $1 - \delta$:*

$$\begin{aligned} \mathcal{S}(\hat{\theta}_c) &\leq O \left(\frac{M \log(n/\delta)}{n} + \mathbb{E}[\text{BIAS}(X; \hat{g})^2] \right) \\ \mathcal{S}(\hat{\theta}_c) &\leq O \left(\sqrt{\frac{\log(M/\delta)}{n}} + \mathbb{E}[|\text{BIAS}(X; \hat{g})|] \right) \end{aligned}$$

As we show in the appendix, the benchmark in the oracle selection guarantee in both the above theorems can be strengthened to be the best model in the convex hull of \mathcal{F} and not just the best in \mathcal{F} . Moreover, we note that the slow rate result of $\sqrt{\log(M/\delta)/n}$ is closely related to the work of Han and Wu (2022), who provided the exact same rate for convex-ERM ensembling using the doubly robust loss that we will introduce in Section 5, albeit only for the case of randomized trials, where we have $\mathbb{E}[|\text{BIAS}(X; \hat{g})|] = 0$, due to knowing the propensity function $\Pr(D = 1 \mid X)$. Thus this result is a direct generalization of their result to observational data. None of these aforementioned theorems achieves the optimal rate when we want to compete with the best in \mathcal{F} . The inability of ERM based approaches to achieve the optimal oracle selection rates motivates the main result of this paper in the next section: the use of ensemble approaches based on Q-aggregation.

4 Q-AGGREGATION WITH NUISANCES

We assume that the function ℓ is σ -strongly convex in f in expectation, i.e., with $\ell^{(2)}$ being the derivative of ℓ with respect to $f(X)$:

$$\begin{aligned} P(\ell(Z; f'(X), g_0) - \ell(Z; f(X), g_0)) &\geq \\ P\ell^{(2)}(Z; f(X), g_0)(f'(X) - f(X)) &+ \frac{\sigma}{2} \|f - f'\|^2 \end{aligned}$$

Consider a modified loss, which penalizes the weight of each model, based on individual model performance:

$$\tilde{\ell}_{\theta,g}(Z) = (1 - \nu)\ell_{\theta,g}(Z) + \nu \sum_{j=1}^M \theta_j \ell_{e_j,g}(Z)$$

For any prior weights $\{\pi_1, \dots, \pi_M\}$, the Q-aggregation ensemble is the solution to the convex problem:

$$\hat{\theta} = \arg \min_{\theta \in \Theta} P_n \tilde{\ell}_{\theta,\hat{g}}(Z) + \frac{\beta}{n} \sum_{j=1}^M \theta_j \log \left(\frac{1}{\pi_j} \right) \quad (\text{Q-agg})$$

Theorem 3 (Main Theorem). *Assume that for some function $\text{ERROR}(\hat{g})$:*

$$\begin{aligned} P(\tilde{\ell}_{\theta,g_0}(Z) - \tilde{\ell}_{\theta',g_0}(Z)) - P(\tilde{\ell}_{\theta,\hat{g}}(Z) - \tilde{\ell}_{\theta',\hat{g}}(Z)) \\ \leq \text{ERROR}(\hat{g}) \|f_{\theta} - f_{\theta'}\|^{\gamma} \end{aligned}$$

for some $\gamma \in [0, 1]$ and $\ell(f, g_0)$ is σ -strongly convex in expectation with respect to f . Moreover, assume that the loss $\ell(Z; f(X), \hat{g})$ is $C_b(\hat{g})$ -Lipschitz with respect to $f(X)$, and both the loss and the candidate functions are uniformly bounded in $[-b, b]$ almost surely. For any $\mu \leq \min\{\nu, 1 - \nu\} \frac{\sigma}{14}$ and $\beta \geq \left\{ \frac{8C_b^2(1-\nu)^2}{\mu}, 8\sqrt{3}bC_b(1-\nu), \frac{2C_b\nu(\nu C_b + \mu b)}{\mu} \right\}$, w.p. $1 - \delta$:

$$\begin{aligned} R(\hat{\theta}, g_0) &\leq \min_{j=1}^M \left(R(e_j, g_0) + \frac{\beta}{n} \log(1/\pi_j) \right) \\ &\quad + O \left(\frac{\log(1/\delta) \max\left\{ \frac{1}{\mu} C_b(\hat{g})^2, b C_b(\hat{g}) \right\}}{n} \right) \\ &\quad + \left(\frac{1}{\mu} \right)^{\frac{\gamma}{2-\gamma}} \text{ERROR}(\hat{g})^{\frac{2}{2-\gamma}} \end{aligned}$$

The theorem gives a more refined version of oracle selection that is adaptive to prior beliefs on which model is best and can even save the $\log(M)$ term, if the prior beliefs end up being correct. In its simplest form, when prior weights are uniform, then β is irrelevant in the optimization problem and can be taken to satisfy the condition of the theorem, as part of the analysis, i.e. the method does not have any hyperparameter β to choose. Moreover, we can always take $\nu = 1/2$, in which case the ensemble has the oracle selection error:

$$S(\hat{\theta}) \leq O \left(\frac{1}{\sigma} \frac{\log(M/\delta)}{n} + \left(\frac{1}{\sigma} \right)^{\frac{\gamma}{2-\gamma}} \text{ERROR}(\hat{g})^{\frac{2}{2-\gamma}} \right)$$

Remark 1 (Computational Considerations). *The computation of the weights θ that optimize the Q-aggregation objective corresponds to a convex optimization problem, which can be solved efficiently by modern convex optimization solvers. Drawing from the results in Dai et al. (2012), we also show in Appendix I that a*

greedy approximate solution to the Q-aggregation problem provides the same statistical guarantees as in Theorem 3. The greedy approximation requires only finding the best single model to the Q objective, as well as the best convex combination between that best single model and any other model (see Algorithm 1). This will always return a convex combination of at most two baseline models and can be implemented by a simple line search over the scalar that represents the convex combination of the two models, as well as the procedure that finds the minimum of the objective over the M models (which is linear in M).

4.1 Square Losses with Nuisance Functions

We specialize the main theorem to the case of square losses, where the target label depends on unknown and estimated nuisance functions as defined in Equation (1), with target parameter τ_0 as defined in Equation (2). For such loss functions, we can instantiate the main theorem to obtain the following corollary.

Corollary 4 (Square Losses with Nuisances). *Let:*

$$\text{ERROR}(\hat{g}) = 2\sqrt{\mathbb{E}[\text{BIAS}(X; \hat{g})^2]}$$

Assume that $\hat{Y}(\hat{g})$ and all functions f_j are absolutely uniformly bounded by U . Then the Q-aggregation ensemble $\hat{\theta}$ with $\nu = 1/2$, $\beta \geq \max\{112U^2, 56U^3\}$ and loss $(\hat{Y}(\hat{g}) - f(X))^2$, satisfies w.p. $1 - \delta$:

$$\begin{aligned} \|f_{\hat{\theta}} - \tau_0\|^2 &\leq \min_{j=1}^M \|f_j - \tau_0\|^2 + \frac{\beta}{n} \log(1/\pi_j) \\ &\quad + O \left(\frac{\log(1/\delta)}{n} + \text{ERROR}(\hat{g})^2 \right) \end{aligned}$$

If we impose more restrictions on the functions in \mathcal{F} , then we can show a corollary, with weaker requirements on the nuisance functions:

Corollary 5. *Assume that for all $\theta, \theta' \in \Theta$:*

$$\|f_{\theta}(X) - f_{\theta'}(X)\|_{L^\infty} \leq C \|f_{\theta}(X) - f_{\theta'}(X)\|_{L^2}^{\gamma}$$

for some constants C and $0 < \gamma \leq 1$. Moreover, let:

$$\text{ERROR}(\hat{g}) = 2C \mathbb{E}[\|\text{BIAS}(X; \hat{g})\|]$$

Then under the remainder conditions and definitions of Corollary 4, w.p. $1 - \delta$

$$\begin{aligned} \|f_{\hat{\theta}} - \tau_0\|^2 &\leq \min_{j=1}^M \|f_j - \tau_0\|^2 + \frac{\beta}{n} \log(1/\pi_j) \\ &\quad + O \left(\frac{\log(1/\delta)}{n} + \text{ERROR}(\hat{g})^{\frac{2}{2-\gamma}} \right) \end{aligned}$$

If the function class \mathcal{F} contains smooth enough functions, then the above condition tends to hold. For

instance, (Mendelson and Neeman, 2010, Lemma 5.1) shows that if the functions $f_j \in F$ lie in an RKHS with a polynomial eigendecay, then the above condition holds for some constant γ that depends on the rate of eigendecay. Moreover, as we show in Appendix F.1, even when \mathcal{F} does not contain smooth functions, we can still show the above property for $\gamma = 1$ and $C \sim M$, as long as functions in \mathcal{F} are not co-linear. Here, the number of models M and not the logarithm of M appears. However, assuming that the nuisance error term decays faster than $n^{-1/2}$, this linear in M term appears only in lower-order terms.

Guarantees without Sample Splitting When learning with nuisance functions, it is often desirable to train the nuisance functions on a separate sample, as it reduces overfitting bias. However, sample splitting decreases efficiency and statistical power, and, thus, it might not be beneficial for small datasets. The following theorem provides theoretical guarantees when the nuisance estimate and the ensemble weights are constructed using the same sample.

Theorem 6 (Guarantees without Sample Splitting). *Under the same conditions as in Corollary 4, with the exception that the nuisance function \hat{g} is trained on the same sample as the parameter $\hat{\theta}$. The Q-aggregation ensemble $\hat{\theta}$ satisfies w.p. $1 - \delta$:*

$$\|f_{\hat{g}} - \tau_0\|^2 \leq \min_j \|f_j - \tau_0\|^2 + \frac{\beta}{n} \log(1/\pi_j) + O\left(\frac{\log(1/\delta)}{n} + r_n^2 + \text{ERROR}(\hat{g})^2\right)$$

where $\text{ERROR}(\hat{g})^2$ is of the order:

$$\mathbb{E} [\text{BIAS}(X; \hat{g})^2] + \mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^4]$$

and r_n denotes the critical radius of the function space $\text{star}(Y_G)$ where $Y_G := \{(\hat{Y}(g) - \hat{Y}(g_0))^2 \mid g \in G\}$.

4.2 Connections to Neyman Orthogonality

In this section, we draw a connection between the bias in labels and the literature of Neyman orthogonality. When we are interested in estimating an average treatment effect, then we can typically phrase this as estimating a parameter of the form: $\tau_0 = \mathbb{E}[\hat{Y}(g_0)]$, for some appropriately defined function $\hat{Y}(g)$. In such settings, we say that the moment function $M(g) := \mathbb{E}[\hat{Y}(g)]$ is Neyman orthogonal (Chernozhukov et al., 2018) with respect to the nuisance functions g if for all $g \in G$, the following directional derivative is zero:

$$\partial_\epsilon M(g_0 + \epsilon(g - g_0)) \big|_{\epsilon=0} = 0$$

Similarly, we can define a conditional analogue, when the quantity of interest is a conditional expectation:

$\tau_0(X) = \mathbb{E}[\hat{Y}(g_0) \mid X]$. We say that the conditional moment function $M_X(g) = \mathbb{E}[\hat{Y}(g) \mid X]$ is conditionally Neyman orthogonal if for all $g \in G$, a.s. over X :

$$\partial_\epsilon M_X(g_0 + \epsilon(g - g_0)) \big|_{\epsilon=0} = 0$$

When the label function $\hat{Y}(g)$ satisfies the latter property then by a second order Taylor expansion we have:

$$\text{BIAS}(X; \hat{g}) = \int_0^1 \partial_\epsilon^2 M_X(g_0 + \epsilon(\hat{g} - g_0)) d\epsilon$$

Thus we can instantiate all the theorems in the aforementioned sections, by using the above bias term. Typically this second order remainder term will depend quadratically in the error of the nuisances, i.e.

$$\int_0^1 \partial_\epsilon^2 M_X(g_0 + \epsilon(\hat{g} - g_0)) d\epsilon \lesssim \mathbb{E} [(\hat{g}(Z) - g_0(Z))^2 \mid X]$$

or will contain only products of errors of the different nuisance functions that compose g (see next section). In both cases, the $\text{ERROR}(\hat{g})^2$ term, will most times be of lower order than $\log(M)/n$ or $\min_j \|f_j - \tau_0\|_{L^2}^2$. The next section and Appendix 6, draw on this connection and provide two Neyman orthogonal label functions for the case of model selection for CATE under unobserved confounding and with access to an instrument.

5 DOUBLY ROBUST Q-AGGREGATION

We now focus on our main question: model selection for CATE, i.e., $\tau_0(X) = \mathbb{E}[Y(1) - Y(0) \mid X]$. We focus on the setting where we observe all confounding variables and we defer to Appendix 6 the case of unobserved confounding with access to an instrumental variable (e.g. A/B testing with non-compliance).

In the no unobserved confounding case, we observe a set of variables W (that is, a superset of X), such that the conditional exogeneity property holds: $Y(1), Y(0) \perp\!\!\!\perp D \mid W$. Under conditional exogeneity, the CATE is identified as:

$$\tau_0(X) = \mathbb{E}[h_0(1, W) - h_0(0, W) \mid X]$$

where $h_0(D, W) := \mathbb{E}[Y \mid D, W]$. As is well known Oprescu et al. (2019); Semenova and Chernozhukov (2021); Foster and Syrgkanis (2023); Kennedy (2020), this identifying formula can be robustified by incorporating propensity estimation: $\tau_0(X) = \mathbb{E}[\hat{Y}(g_0) \mid X]$. This gives rise to the doubly robust square loss:

$$\ell(Z; f(X), g) = (\hat{Y}(g) - f(X))^2 \quad (\text{DR-Loss})$$

For any estimate $\hat{g} = (\hat{h}, \hat{a})$, define $\hat{g}(D, W)$ as:

$$(a_0(D, W) - \hat{a}(D, W))(h_0(D, W) - \hat{h}(D, W))$$

An important fact of the doubly robust target, exploited also in prior works, is that it satisfies the mixed bias property (see Appendix C):

$$\text{BIAS}(X; \hat{g}) := \mathbb{E} \left[\hat{Y}(g_0) - \hat{Y}(\hat{g}) \mid X \right] = \mathbb{E} [\hat{q}(D, W) \mid X]$$

Thus we can apply Corollary 4 to obtain:

Corollary 7 (Main Corollary). *Let $\tau_0(X) = \mathbb{E}[Y(1) - Y(0) \mid X]$ and suppose that conditional exogeneity holds conditional on some observed W that is a superset of X . Let:*

$$\text{ERROR}(\hat{g}) = 2\sqrt{\mathbb{E} \left[\mathbb{E} [\hat{q}(D, W) \mid X]^2 \right]}$$

and assume that $\hat{Y}(\hat{g})$ all functions $f \in F$ are uniformly and absolutely bounded by U . Then the doubly robust Q-aggregation ensemble $\hat{\theta}$, based on the DR-Loss, with $\nu = 1/2$ and a uniform prior π , satisfies that w.p. $1 - \delta$:

$$\begin{aligned} \|f_{\hat{\theta}} - \tau_0\|^2 &\leq \min_j \|f_j - \tau_0\|^2 \\ &\quad + O \left(\frac{\log(M/\delta)}{n} + \text{ERROR}(\hat{g})^2 \right) \end{aligned}$$

For any prior π , if $\beta \geq \max\{112U^2, 56U^3\}$, w.p. $1 - \delta$:

$$\begin{aligned} \|f_{\hat{\theta}} - \tau_0\|^2 &\leq \min_j \|f_j - \tau_0\|^2 + \frac{\beta}{n} \log(1/\pi_j) \\ &\quad + O \left(\frac{\log(1/\delta)}{n} + \text{ERROR}(\hat{g})^2 \right) \end{aligned}$$

The error term can also be further upper bounded by an application of a Cauchy-Schwarz inequality by:

$$\begin{aligned} \text{ERROR}(\hat{g}) &\leq \min\{\|\hat{a} - a_0\|_{L^4} \|\hat{h} - h_0\|_{L^4}, \\ &\quad \|\hat{a} - a_0\|_{L^2} \|\hat{h} - h_0\|_{L^\infty}, \|\hat{a} - a_0\|_{L^\infty} \|\hat{h} - h_0\|_{L^2}\} \end{aligned}$$

Thus we can either control the L^4 moment of the prediction error of the two nuisances, or guarantee an L^∞ rate for one and an L^2 rate for the other. Note that if the above product of error terms is of lower order than $n^{-1/2}$, then the effect of these errors on the quality of the causal ensemble is of second-order importance.

If we further make an assumption that:

$$\|f_{\theta^*}(X) - f_{\hat{\theta}}(X)\|_{L^\infty} \leq C \|f_{\theta^*}(X) - f_{\hat{\theta}}(X)\|_{L^2}^\gamma$$

Then by Corollary 4, we can take in Corollary 7:

$$\text{ERROR}(\hat{g}) \lesssim C \|\hat{a} - a_0\|_{L^2}^{\frac{1}{2-\gamma}} \|\hat{h} - h_0\|_{L^2}^{\frac{1}{2-\gamma}}$$

As we saw in Section 4.1, the latter assumption holds if either the functions in \mathcal{F} are smooth or they contain independent components in their predictions.

Remark 2 (Priors on CATE models). *In the context of CATE estimation, incorporating priors is quite natural. For instance, we have a strong prior that models that come out of meta-learning approaches that only use outcome modeling should perform worse. Thus, we can put lower prior probability on S- and T- learner models and larger prior probability on X-, DR- and R- learner models. Moreover, we can incorporate priors that more complicated models are less probable than simpler models. For instance, we can put higher prior on linear cate models or shallow random forest models.*

6 DOUBLY ROBUST Q-AGGREGATION FOR CATE WITH INSTRUMENTS

We consider a widely encountered setting of estimating a CATE from a stratified randomized trial with non-compliance. In this setting we know the cohort assignment policy $\pi_0(X) = \Pr(Z = 1 \mid X)$, and we are interested in estimating the conditional local average treatment effect

$$\tau_0(X) = \mathbb{E}[Y(1) - Y(0) \mid D(1) > D(0), X]$$

where $D(z)$ is the potential outcome of the chosen treatment under different cohort assignments (aka recommended treatments). This effect is identified as:

$$\begin{aligned} \tau_0(X) &= \frac{\mathbb{E}[Y\tilde{Z} \mid X]}{\mathbb{E}[D\tilde{Z} \mid X]} = \frac{\text{Cov}(Y, Z \mid X)}{\text{Cov}(D, Z \mid X)} \\ &= \frac{\mathbb{E}[Y \mid Z = 1, X] - \mathbb{E}[Y \mid Z = 0, X]}{\mathbb{E}[D \mid Z = 1, X] - \mathbb{E}[D \mid Z = 0, X]} \end{aligned}$$

where $\tilde{Z} = Z - \pi_0(X)$.

Let \hat{a} be an estimate of:

$$\begin{aligned} a_0(X) &:= \mathbb{E}[Y\tilde{Z} \mid X] \\ &= \mathbb{E}[Y \mid Z = 1, X] - \mathbb{E}[Y \mid Z = 0, X] \end{aligned}$$

and $\hat{\beta}$ an estimate of:

$$\begin{aligned} \beta_0(X) &:= \mathbb{E}[D\tilde{Z} \mid X] \\ &= \mathbb{E}[D \mid Z = 1, X] - \mathbb{E}[D \mid Z = 0, X] \end{aligned}$$

and let $\hat{\tau} = \hat{a}/\hat{\beta}$. Then we can construct the random variable

$$\hat{Y}(\hat{g}) = \hat{\tau}(X) + \frac{(Y - \hat{\tau}(X)D)\tilde{Z}}{\hat{\beta}(X)}$$

and we can consider Q-aggregation with the loss:

$$\ell(Z; f(X), g) = (\hat{Y}(g) - f(X))^2$$

In this case, we can show the following mixed bias property (see Appendix C):

$$\mathbb{E}[\hat{Y}(\hat{g}) - \hat{Y}(g_0) | X] = (\tau_0(X) - \hat{\tau}(X)) \frac{\beta_0(X) - \hat{\beta}(X)}{\hat{\beta}(X)}$$

Thus applying Corollary 4 we have:

Corollary 8 (Main Corollary for CATE with Instruments). *Let $\tau_0(X) = \mathbb{E}[Y(1) - Y(0) | D(1) > D(0), X]$ and suppose that the instrument is a stratified randomized trial with non-compliance and with a known propensity model $\pi_0(X) = \Pr(Z = 1 | X)$. Let:*

$$\text{ERROR}(\hat{g})^2 = 4\mathbb{E} \left[(\tau_0(X) - \hat{\tau}(X))^2 \frac{(\beta_0(X) - \hat{\beta}(X))^2}{\hat{\beta}(X)^2} \right]$$

and assume that $\hat{Y}(\hat{g})$ is uniformly and absolutely bounded by U in a sufficiently small neighborhood of g_0 and that all functions $f \in F$ are uniformly and absolutely bounded by U . Then the doubly robust IV Q-aggregation ensemble $\hat{\theta}$ with $\nu = 1/2$ and a uniform prior π , satisfies that w.p. $1 - \delta$:

$$\begin{aligned} \|f_{\hat{\theta}} - \tau_0\|^2 &\leq \min_j \|f_j - \tau_0\|^2 \\ &+ O \left(\frac{\log(M) + \log(1/\delta)}{n} + \text{ERROR}(\hat{g})^2 \right) \end{aligned}$$

Moreover, for any prior π , for $\beta \geq \max\{112U^2, 56U^3\}$, w.p. $1 - \delta$:

$$\begin{aligned} \|f_{\hat{\theta}} - \tau_0\|^2 &\leq \min_j \|f_j - \tau_0\|^2 + \frac{\beta}{n} \log(1/\pi_j) \\ &+ O \left(\frac{\log(1/\delta)}{n} + \text{ERROR}(\hat{g})^2 \right) \end{aligned}$$

Moreover, note that since $\hat{\tau} = \hat{\alpha}/\hat{\beta}$, we also have that

$$\tau_0 - \hat{\tau} = \frac{\alpha_0\hat{\beta} - \hat{\alpha}\beta_0}{\hat{\beta}\beta_0} = \frac{(\alpha_0 - \hat{\alpha})\hat{\beta} + \hat{\alpha}(\hat{\beta} - \beta_0)}{\hat{\beta}\beta_0}$$

Thus overall we have:

$$\begin{aligned} \text{ERROR}(\hat{g})^2 &= 4\mathbb{E} \left[\frac{(\alpha_0(X) - \hat{\alpha}(X))^2 (\beta_0(X) - \hat{\beta}(X))^2}{\beta_0(X)^2 \hat{\beta}(X)^2} \right] \\ &+ 4\mathbb{E} \left[\frac{\hat{\alpha}(X)^2 (\beta_0(X) - \hat{\beta}(X))^4}{\beta_0(X)^2 \hat{\beta}(X)^4} \right] \end{aligned}$$

Assuming that $\beta_0(X) \geq \underline{\beta}$ and that we impose the same restriction on our estimate $\hat{\beta}(X)$ (which is a "minimal conditional instrument strength") and assuming that $\hat{\alpha}(X)^2 \leq \bar{\alpha}^2$, then we have:

$$\begin{aligned} \text{ERROR}(\hat{g})^2 &\lesssim \mathbb{E} \left[(\alpha_0(X) - \hat{\alpha}(X))^2 (\beta_0(X) - \hat{\beta}(X))^2 \right] \\ &+ \mathbb{E} \left[(\beta_0(X) - \hat{\beta}(X))^4 \right] \end{aligned}$$

By Cauchy-Schwarz inequality the latter is upper bounded by:

$$\begin{aligned} \min\{ \|\hat{\alpha} - \alpha_0\|_{L^4}^2 \|\hat{\beta} - \beta_0\|_{L^4}^2 + \|\hat{\beta} - \beta_0\|_{L^4}^2, \\ \|\hat{\alpha} - \alpha_0\|_{L^2}^2 \|\hat{\beta} - \beta_0\|_{L^\infty}^2 + \|\hat{\beta} - \beta_0\|_{L^\infty}^2 \} \end{aligned}$$

Thus if we learn the effect of the instrument on the treatment (typically referred to as the compliance model) with an L^∞ guarantee (e.g. a high-dimensional sparse linear logistic regression), then it suffices to learn the effect of the instrument on the outcome with an L^2 error. Moreover, as long as the error on the compliance model $\hat{\beta}$ is $o(n^{-1/4})$ and the product of the errors of $\hat{\alpha}$ and $\hat{\beta}$ is $o(n^{-1/2})$, then the nuisance impact is of lower order in the final guarantee for the Q-aggregation ensemble.

Again, if we further make an assumption that:

$$\|f_{\theta^*}(X) - f_{\hat{\theta}}(X)\|_{L^\infty} \leq C \|f_{\theta^*}(X) - f_{\hat{\theta}}(X)\|_{L^2}$$

Then we can take:

$$\begin{aligned} \text{ERROR}(\hat{g}) &= 2C \mathbb{E} \left[\left| \mathbb{E}[\hat{Y}(\hat{g}) - \hat{Y}(g_0) | X] \right| \right] \\ &\lesssim \|\hat{\alpha} - \alpha_0\|_{L^2} \|\hat{\beta} - \beta_0\|_{L^2} + \|\hat{\beta} - \beta_0\|_{L^2}^2 \end{aligned}$$

and it suffices to control only the RMSE errors of $\hat{\alpha}$ and $\hat{\beta}$ to be $o(n^{-1/4})$.

7 EXPERIMENTS

We assessed the performance of our doubly robust Q-aggregation method on fully simulated and semi-synthetic data sets. We constructed several candidate CATE models based on meta-learning approaches using XGboost for regression and classification subproblems, such as for estimating the nuisance functions: $\hat{\mu}(X, D)$ (an estimate of $\mathbb{E}[Y | X, D]$), $\hat{\mu}_d(X)$ (an estimate of $\mathbb{E}[Y | X, D = d]$), and $\hat{\pi}(X)$ (an estimate for the propensity). For constructing the CATE models we considered 8 meta-learning strategies, namely S-, T-, IPW-, X-, DR-, R-, DRX-, and DAX-learners. Details of each learner are given in Appendix D.1.

We present results for Q-aggregation, convex stacking, and best-ERM (selecting the model with the best doubly robust loss). Note that convex stacking is a special case of Q-aggregation with ν set to 0. For Q-aggregation, we chose $\nu = 0.1$ and did not incorporate any priors. Each data set is divided into 3 portions: 60%, 20%, and 20%. Candidate CATE models and nuisance functions are trained on the 60% of the data. Subsequently, the CATE ensemble weights are trained on 20% of the data. Finally, CATE RMSE is evaluated on the remaining 20% of the data. For each model, we report the RMSE regret when compared to the oracle model selection - the model that achieves the lowest RMSE on the test set (i.e. $\|f_{\hat{\theta}} - \tau_0\|^2 - \min_j \|f_j - \tau_0\|^2$).

Table 1: Mean RMSE Regret Over DGPs And Semi-synthetic Datasets.

	S	T	IPS	X	DR	R	DRX	DAX	Qtrain	Convextrain	Besttrain
DGP	1.495	0.946	2.350	1.149	0.663	1.125	0.894	0.837	0.485	0.488	0.567
Simple Semi-synthetic	0.611	0.534	7.081	0.284	0.445	0.619	0.411	0.384	0.194	0.204	0.234
Fitted Semi-Synthetic	0.133	1.912	5.106	0.431	0.417	0.899	0.407	0.816	0.300	0.322	0.258
All Experiments	0.746	1.131	4.846	0.621	0.508	0.881	0.570	0.679	0.327	0.338	0.353

Table 2: RMSE Regret For Semi-synthetic Datasets with Simple CATE.

	401k	welfare	poverty	star	criteo
Besttrain	[355 ± 147.9] 362.4 (585.7)	[0.0390 ± 0.0181] 0.0406 (0.0722)	[0.0075 ± 0.0068] 0.0060 (0.0194)	[0.0107 ± 0.0096] 0.0070 (0.0287)	[0.00081 ± 0.00040] 0.0008 (0.00161)
Convextrain	[362.3 ± 157.1] 371.5 (634.4)	[0.0289 ± 0.0156] 0.0275 (0.0557)	[0.0066 ± 0.0049] 0.0057 (0.0166)	[0.0095 ± 0.0093] 0.0062 (0.0284)	[0.00066 ± 0.00032] 0.0006 (0.00122)
Qtrain	[341.1 ± 157.1] 331.3 (634.2)	[0.0281 ± 0.0156] 0.0272 (0.0537)	[0.0061 ± 0.0049] 0.0051 (0.0165)	[0.0095 ± 0.0093] 0.0061 (0.0284)	[0.00060 ± 0.00032] 0.0006 (0.00112)
% Decrease w.r.t Best	4.0692 % 9.4008 %	38.7569 % 49.2038 %	22.8005 % 17.7918 %	13.0432 % 15.7684 %	34.8879 % 30.0913 %
% Decrease w.r.t Convex	6.2148 % 12.1341 %	2.6238 % 1.0101 %	8.1161 % 13.0241 %	0.2770 % 2.3870 %	9.9629 % 8.5184 %

Simulated Data We employed six distinct data generation processes and performed experiments (DGP) on 100 instances of each DGP. In all cases, a scalar covariate X is drawn from a uniform distribution. Each DGP is characterized by 3 functions: propensity $\pi(X)$, conditional baseline response surface $m(X)$, and treatment effect $\tau(X)$. Treatment D is sampled from a binomial distribution with probability $\pi(X)$, and outcome Y is simulated by:

$$Y = m(X) + D\tau(X) + \epsilon$$

where $\epsilon \sim N(0, \sigma^2)$. The details of each DGP can be found in the supplementary materials.

Semi-synthetic Data Semi-synthetic data was generated from five datasets used in previous studies: welfare dataset (Green and Kern, 2012), poverty dataset (Farbmacher et al., 2021), Project STAR dataset (Word et al., 1990), 401k eligibility dataset (Poterba and Venti, 1994; Poterba et al., 1995), and the Criteo Uplift Modeling (Diemert Eustache, Betlei Artem et al., 2018) dataset. Since our analysis assumes a binary treatment, only one treatment is considered for datasets with multiple treatments.

We used two simulation modes to generate the outcome from covariates. In the simple mode, the outcome is generated based on a simple linear CATE and baseline response function plus Gaussian noise (with $\sigma = 0.2$). In the fitted mode, the output is generated from a random forest regressor that is fitted on the real dataset, with randomness introduced by the addition of the model residual uniformly at random.

Results Overall, there is no single meta-learning model that performs consistently well across all DGPs

or semi-synthetic datasets. We present in Table 1 the mean RMSE regret for different DGPs and different semisynthetic datasets, where the RMSE regret is normalized by the mean RMSE regret of all models for each DGP/dataset. The results in Table 1 show that best-model selection, convex stacking, and Q aggregation achieve a significantly lower average RMSE regret compared to choosing any single model.

On semi-synthetic datasets, Q-aggregation achieves comparable RMSE regret with convex stacking (see Figure 2) and consistently outperforms both best-ERM and convex stacking in the case of a simple CATE and baseline response function. Table 2 summarizes the RMSE regret for model selection, convex stacking, and Q-aggregation for semi-synthetic datasets with simple CATE functions, where each cell in the table reports [mean ± st.dev.] median (95%) of RMSE Regret over 100 experiments for each semi-synthetic dataset simulated using the simple mode. The last two rows present the percentage decrease in the mean and median RMSE regret for Q-aggregation with respect to convex stacking and model selection.

For simple CATE functions, Table 2 shows that there is an decrease in the mean and median RMSE regret for Q-aggregation, when compared with convex stacking and model selection for every dataset. For datasets generated with the fitted CATE functions, the underlying model follows a S-learner setup. Thus, we expect the S-learners to be dominant models, which is evident Table 1. As a result, we observe that there is often an increase in the RMSE regret for ensemble models over the best single model, as is hinted by Theorem 1, since in this case we have a dominant model.

References

- Alaa, A. and Van Der Schaar, M. (2019). Validating causal inference models via influence functions. In *International Conference on Machine Learning*, pages 191–201. PMLR.
- Athey, S., Bayati, M., Imbens, G., and Qu, Z. (2019). Ensemble methods for causal effects in panel data settings. In *AEA papers and proceedings*, volume 109, pages 65–70. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- Battocchi, K., Dillon, E., Hei, M., Lewis, G., Oka, P., Oprescu, M., and Syrgkanis, V. (2019). EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation. <https://github.com/py-why/EconML>. Version 0.14.1.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters.
- Chernozhukov, V., Newey, W., Quintas-Martinez, V. M., and Syrgkanis, V. (2022). Riesznet and forests: Automatic debiased machine learning with neural nets and random forests. In *International Conference on Machine Learning*, pages 3901–3914. PMLR.
- Curth, A. and van der Schaar, M. (2023). In search of insights, not magic bullets: Towards demystification of the model selection dilemma in heterogeneous treatment effect estimation. *arXiv preprint arXiv:2302.02923*.
- Dai, D., Rigollet, P., and Zhang, T. (2012). Deviation optimal learning using greedy q-aggregation. *The Annals of Statistics*, pages 1878–1905.
- Diemert Eustache, Betlei Artem, Renaudin, C., and Massih-Reza, A. (2018). A large scale benchmark for uplift modeling. In *Proceedings of the AdKDD and TargetAd Workshop, KDD, London, United Kingdom, August, 20, 2018*. ACM.
- Dorie, V., Hill, J., Shalit, U., Scott, M., and Cervone, D. (2019). Automated versus do-it-yourself methods for causal inference. *Statistical Science*, 34(1):43–68.
- Dwivedi, R., Tan, Y. S., Park, B., Wei, M., Horgan, K., Madigan, D., and Yu, B. (2020). Stable discovery of interpretable subgroups via calibration in causal studies. *International Statistical Review*, 88:S135–S178.
- Farbmacher, H., Kögel, H., and Spindler, M. (2021). Heterogeneous effects of poverty on attention. *Labour Economics*, 71:102028.
- Foster, D. J. and Syrgkanis, V. (2023). Orthogonal statistical learning. *The Annals of Statistics*, 51(3):879–908.
- Green, D. P. and Kern, H. L. (2012). Modeling heterogeneous treatment effects in survey experiments with bayesian additive regression trees. *Public opinion quarterly*, 76(3):491–511.
- Grimmer, J., Messing, S., and Westwood, S. J. (2017). Estimating heterogeneous treatment effects and the effects of heterogeneous treatments with ensemble methods. *Political Analysis*, 25(4):413–434.
- Hahn, P. R., Murray, J. S., and Carvalho, C. M. (2020). Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Analysis*, 15(3):965–1056.
- Han, K. W. and Wu, H. (2022). Ensemble method for estimating individualized treatment effects. *arXiv preprint arXiv:2202.12445*.
- Hartford, J., Lewis, G., Leyton-Brown, K., and Taddy, M. (2017). Deep IV: A flexible approach for counterfactual prediction. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1414–1423. PMLR.
- Imai, K. and Ratkovic, M. (2013). Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics*, 7(1):443 – 470.
- Kennedy, E. H. (2020). Towards optimal doubly robust estimation of heterogeneous causal effects. *arXiv preprint arXiv:2004.14497*.
- Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10):4156–4165.
- Lecué, G. and Rigollet, P. (2014). Optimal learning with Q-aggregation. *The Annals of Statistics*, 42(1):211 – 224.
- Ledoux, M. and Talagrand, M. (1991). *Probability in Banach Spaces: isoperimetry and processes*, volume 23. Springer Science & Business Media.
- Mahajan, D., Mitliagkas, I., Neal, B., and Syrgkanis, V. (2022). Empirical analysis of model selection for heterogeneous causal effect estimation. *arXiv preprint arXiv:2211.01939*.
- McAlinn, K. and West, M. (2019). Dynamic bayesian predictive synthesis in time series forecasting. *Journal of econometrics*, 210(1):155–169.

- Mendelson, S. and Neeman, J. (2010). Regularization in kernel learning. *The Annals of Statistics*, pages 526–565.
- Mitchell, C. and van de Geer, S. (2009). General oracle inequalities for model selection. *Electronic Journal of Statistics*, 3:176–204.
- Nie, X. and Wager, S. (2021). Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319.
- Oprescu, M., Syrgkanis, V., and Wu, Z. S. (2019). Orthogonal random forest for causal inference. In *International Conference on Machine Learning*, pages 4932–4941. PMLR.
- Poterba, J. M. and Venti, S. F. (1994). 401 (k) plans and tax-deferred saving. In *Studies in the Economics of Aging*, pages 105–142. University of Chicago Press.
- Poterba, J. M., Venti, S. F., and Wise, D. A. (1995). Do 401 (k) contributions crowd out other personal saving? *Journal of Public Economics*, 58(1):1–32.
- Rolling, C. A. and Yang, Y. (2014). Model selection for estimating treatment effects. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 76(4):749–769.
- Saito, Y. and Yasui, S. (2020). Counterfactual cross-validation: Stable model selection procedure for causal inference models. In *International Conference on Machine Learning*, pages 8398–8407. PMLR.
- Schuler, A., Baiocchi, M., Tibshirani, R., and Shah, N. (2018). A comparison of methods for model selection when estimating individual treatment effects. *arXiv preprint arXiv:1804.05146*.
- Semenova, V. and Chernozhukov, V. (2021). Debiased machine learning of conditional average treatment effects and other causal functions. *The Econometrics Journal*, 24(2):264–289.
- Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*. Cambridge university press.
- Shalit, U., Johansson, F. D., and Sontag, D. (2017). Estimating individual treatment effect: generalization bounds and algorithms. In *International conference on machine learning*, pages 3076–3085. PMLR.
- Shi, C., Blei, D., and Veitch, V. (2019). Adapting neural networks for the estimation of treatment effects. *Advances in neural information processing systems*, 32.
- Sugasawa, S., Takanashi, K., and McAlinn, K. (2023). Bayesian causal synthesis for supra-inference on heterogeneous treatment effects. *arXiv preprint arXiv:2304.07726*.
- Syrgkanis, V., Lei, V., Oprescu, M., Hei, M., Battocchi, K., and Lewis, G. (2019). Machine learning estimation of heterogeneous treatment effects with instruments. *Advances in Neural Information Processing Systems*, 32.
- Tsybakov, A. B. (2003). Optimal rates of aggregation. In *Learning Theory and Kernel Machines: 16th Annual Conference on Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24-27, 2003. Proceedings*, pages 303–313. Springer.
- Vaart, A. W. v. d., Dudoit, S., and Laan, M. J. v. d. (2006). Oracle inequalities for multi-fold cross validation. *Statistics & Decisions*, 24(3):351–371.
- Van der Laan, M., Polley, E., and Hubbard, A. (2007). Super learner, statistical applications in genetics and molecular biology. *Super learner. Statistical applications in genetics and molecular biology*, 6(1).
- Van Der Laan, M. J. and Dudoit, S. (2003). Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples.
- Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.
- Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press.
- Word, E. et al. (1990). Student/teacher achievement ratio (star) tennessee’s k-3 class size study. final summary report 1985-1990.

Checklist

- For all models and algorithms presented, check if you include:
 - A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes] See supplementary materials.
- For any theoretical claim, check if you include:
 - Statements of the full set of assumptions of all theoretical results. [Yes]

- (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes] See code in the Supplementary materials.
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes] See the Experiments Section.
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes] Standard deviations presented in the tables were obtained from 100 experiments over different random seeds.
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes] The experiments were run on CPUs.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes] See code in the supplementary materials.
 - (d) Information about consent from data providers/curators. [Yes] The datasets used in this study are all public datasets.
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Causal Q-Aggregation for CATE Model Selection: Appendix

Contents

1	INTRODUCTION	1
2	PRELIMINARIES	3
3	ERM-BASED SELECTION	3
4	Q-AGGREGATION WITH NUISANCES	4
4.1	Square Losses with Nuisance Functions	5
4.2	Connections to Neyman Orthogonality	6
5	DOUBLY ROBUST Q-AGGREGATION	6
6	DOUBLY ROBUST Q-AGGREGATION FOR CATE WITH INSTRUMENTS	7
7	EXPERIMENTS	8
A	DETAILED COMPARISONS WITH PRIOR WORKS	15
A.1	COMPARISON TO Van Der Laan and Dudoit (2003)	16
B	PROOFS FOR ERM-BASED SELECTION	18
B.1	Preliminaries	18
B.2	Best-ERM: Proof of Theorem 1 of the Main Paper	21
B.3	Convex-ERM: Proof of Theorem 2 of the Main Paper	22
C	MIXED BIAS PROPERTIES	23
C.1	Doubly Robust Targets for CATE under Conditional Exogeneity	23
C.2	Doubly Robust Targets for CATE with Instruments	23
D	EXPERIMENT DETAILS	24
D.1	Candidate Models	24
D.2	Data Generating Processes for Synthetic Datasets	25
D.3	Semi-synthetic Datasets	25
D.4	Additional Results	25

E	PROOF OF MAIN THEOREM	28
E.1	Concentration of offset empirical processes	32
E.2	Putting it all together	36
E.3	Basic concentration proofs	36
E.4	Auxiliary Lemmas	37
F	PROOF OF COROLLARIES FOR SQUARE LOSSES	37
F.1	Relaxed Nuisance Guarantees Under Eigenvalue Condition	39
G	PROOF OF MAIN THEOREM WITHOUT SAMPLE SPLITTING	39
G.1	Putting it all together	44
H	ADDITIONAL REMARKS FOR CATE WITH INSTRUMENTS	44
I	GREEDY Q-AGGREGATION FOR SQUARE LOSSES	44
I.1	Preliminaries	44
I.2	Excess Risk Guarantees for the k-step Greedy Q-aggregation Estimator	46
J	PROOF OF ORTHOGONAL LEARNING THEOREM	50
J.1	Fast Rates	50
J.2	Slow Rates	52
J.2.1	With Universal Orthogonality	52
J.2.2	Without Universal Orthogonality	52

A DETAILED COMPARISONS WITH PRIOR WORKS

In traditional supervised learning, the standard approach for model selection is through M-fold cross validation, also called the discrete super learner. Vaart et al. (2006) proved an oracle inequality (under conditions):

$$\mathbb{E}[R(\hat{f})] - \mathbb{E}[R(f^*)] \leq O\left(\mathbb{E}[R(f^*)] + \frac{\log(M)}{n}\right)$$

where \hat{f} is the best model selected by K-fold cross validation, f^* is the oracle model selection, and $R(f)$ is the excess risk. Similar oracle inequalities are also provided in Van der Laan et al. (2007). Moreover, Van Der Laan and Dudoit (2003) proved similar oracle inequalities for cross-validation even in the presence of nuisance functions in the loss. Our results on ERM-based selection and Q-aggregation based selection improve upon the results of Van Der Laan and Dudoit (2003) on many critical directions. In summary, the results in Van Der Laan and Dudoit (2003) cannot yield the doubly robust form of the dependence in the errors of the nuisances, for our main application to CATE model selection. Moreover, the selection rate results in Van Der Laan and Dudoit (2003) depend on the performance of the best model in the set and only provide slow rates if the loss of the best model is not close to zero. Given the subtleties of the comparison of the results in Van Der Laan and Dudoit (2003) and our work, we give very detailed comparison in the next subsection.

A similar oracle inequality was shown when ERM is used for model selection (Mitchell and van de Geer, 2009). In addition, under strong assumptions on the tails of the centered excess losses, the authors proved that model selection via ERM can achieve a tail bound that with probability $1 - \delta$:

$$R(\hat{f}) - R(f^*) \leq O\left(R(f^*) + \frac{\log(M)}{n} + \frac{\log(1/\delta)}{n}\right)$$

Model Selection is particularly difficult for CATE estimation due to the lack of ground truth counterfactual data, making it difficult to evaluate model performance. Existing model selection criterion include the influence corrected loss (Alaa and Van Der Schaar, 2019), R-loss for double machine learning (Nie and Wager, 2021), doubly robust loss (Saito and Yasui, 2020), the X-score (Künzel et al., 2019), matching score, (Rolling and Yang, 2014), calibration score (Dwivedi et al., 2020), etc. In two recently conducted benchmarking studies on model selection for CATE estimators (Mahajan et al., 2022; Curth and van der Schaar, 2023), the results showed that no single evaluation metric stood out as the dominant choice. Moreover, Mahajan et al. (2022) demonstrated that softmax stacking of the estimators often lead to better performance for any proxy evaluation metric.

Another line of work employs stacking which optimizes model weights by minimizing some performance metric on a held-out sample. Grimmer et al. (2017) constructed a weighted average across different potential outcome models via a cross validation based super learning approach (Van der Laan et al., 2007) that estimates the weights from out-of-sample performance through regression. In contrast to our work, Grimmer et al. (2017) applies stacking to the nuisance function that is used to construct the estimator, instead of the estimator itself. Closer to our line of work, Nie and Wager (2021) proposed to use a positive linear combination of different estimators that minimizes their proposed R-loss on a validation sample. Although both work demonstrated empirical success, they did not provide any theoretical guarantees for stacking.

Bridging this gap, Han and Wu (2022) studied the the theoretical guarantees of their proposed stacking algorithm when compared to oracle model selection. The proposed causal stacking algorithm learns a convex combination of the candidate models by optimizing the doubly robust loss on a validation set. Notably, what we are presenting in this work, that is Q aggregation, can be reduced to causal stacking when choose the same loss and some hyperparameter is set to 0. Han and Wu showed that, with probability $1 - \delta$, the RMSE regret is on the order of $\sqrt{\frac{\log(M^2)/\delta}{n}}$, which is a slower rate that what we presented in this paper. Interestingly, allowing the weights to be any non-negative value instead of summing to 1 improved performance when the treatment assignment is balanced.

More recently, Sugasawa et al. (2023) sought a different approach to construct ensembles which uses the Bayesian predictive synthesis (McAlinn and West, 2019) framework, which allows the model/posterior weights to depend on the covariates, enabling the final ensemble to capture the heterogeneity of the performance of different models.

In our work, we provide additional theoretical results utilizing the orthogonal statistical learning framework (Foster and Syrgkanis, 2023) when a Neyman orthogonal loss (e.g. the doubly robust loss) is used. In Section

3, we present results for both model selection and convex stacking using plug-in empirical risk minimization. In addition, we propose Q-aggregation (Lecué and Rigollet, 2014) for CATE estimation, which constructs a convex combination of the candidate functions via a modified version of empirical risk minimization, and achieves optimal rate (Tsybakov, 2003) of $O(\frac{\log(M)}{n})$ plus higher order error terms for the nuisance functions.

A.1 COMPARISON TO Van Der Laan and Dudoit (2003)

Van Der Laan and Dudoit (2003) proved oracle inequalities for cross-validation that considers the need to estimate nuisance functions for the loss function. Here, we reproduce the main theorem in Van Der Laan and Dudoit (2003) using our notation.

Theorem 9 (Theorem 1 of Van Der Laan and Dudoit (2003)). *Let $f_{\hat{k}}$ and $f_{\bar{k}}$ denote the best model selected by cross-validation and oracle model selection respectively. Define the excess risk with respect to the true candidate model as $d(f, f_0) = \mathbb{E}_{S_n}[P(\ell_{f,g_0} - \ell_{f_0,g_0})]$, where S_n is the data split. Let p denote the proportion of data used for validation in each split. Further define:*

$$\begin{aligned} r_1(n) &= \max_{\bar{k} \in \{\hat{k}, \bar{k}\}} \frac{\mathbb{E}[P(\ell_{f_{\bar{k}}, \hat{g}} - \ell_{f_0, \hat{g}}) - P(\ell_{f_{\bar{k}}, g_0} - \ell_{f_0, g_0})]}{\sqrt{\mathbb{E}[P(\ell_{f_{\bar{k}}, g_0} - \ell_{f_0, g_0})]}} \\ r_2(n) &= \mathbb{E} \left[\max_{\bar{k}} \sqrt{P((\ell_{f_{\bar{k}}, \hat{g}} - \ell_{f_0, \hat{g}}) - (\ell_{f_{\bar{k}}, g_0} - \ell_{f_0, g_0}))^2} \right] \\ \tilde{r}(n) &= \sqrt{\mathbb{E}[d(f_{\bar{k}}, f_0)]} \end{aligned}$$

Assuming that the nuisance estimator \hat{g} converges to the true nuisance function g_0 as $n \rightarrow \infty$, the loss function is uniformly bounded by some constant M_1 for any $f \in F$, and $\text{Var}(\ell_{f,g_0} - \ell_{f_0,g_0}) \leq M_2 \mathbb{E}[\ell_{f,g_0} - \ell_{f_0,g_0}]$ for some constant M_2 , then for any $\delta > 0$:

$$\sqrt{\mathbb{E}[d(f_{\bar{k}}, f_0)]} \leq \frac{(1 + \delta)r_1(n) + \sqrt{(1 + \delta)^2 r_1(n)^2 + 4\epsilon_n(\delta)}}{2} =: \sqrt{R(n)}$$

where

$$\begin{aligned} \epsilon_n(\delta) &= (1 + 2\delta)\tilde{r}^2(n) + 4(1 + \delta)^2 \left(\frac{2M_1}{3} + \frac{M_2}{\delta} \right) \frac{1 + \log(M)}{np} + \\ &\quad (1 + \delta)r_1(n)\tilde{r}(n) + \frac{2M_3(1 + \delta)\log(M)}{(np)^{0.5}} \max(r_2(n), (np)^{-0.5} I(r_2(n) > 0)) \end{aligned}$$

And M_3 is given by:

$$M_3(n) = \frac{M_1}{3} + \sqrt{2} \frac{\sqrt{\log(2)}}{\log(M)} + \frac{\sqrt{2}}{\sqrt{\log(M)}} + b_0 + \int_{b_0}^{\infty} 2M^{1-m(x)} dx$$

where

$$m(x) = \frac{1}{2} \frac{x^2}{1/\log(M) + \frac{4}{3M_1}x}$$

and b_0 is the solution to $1 - m(x) = 0$.¹

Note that $1 - m(x)$ is decreasing in x . Thus $M_3(n)$ is upper and lower bounded by some strictly positive constant factor that only depends on M_1 . Up to a constant factor $C(M_1, M_2)$ that depends only on M_1, M_2 (and assuming

¹The theorem of Van Der Laan and Dudoit (2003) also showed a form of asymptotic optimality only under the extra condition that:

$$O(r_1(n)^2 + \log(M)r_2(n)^2) = O(\tilde{r}(n)^2)$$

for simplification that $\delta \leq 1$), the above can be simplified as:

$$\begin{aligned} R(n) &\sim (1 + \delta)^2 r_1(n)^2 + 4\epsilon_n(\delta) \\ &\sim (1 + 2\delta)\mathbb{E}[d(f_{\hat{k}}, f_0)] + C(M_1, M_2) \left(\frac{1}{\delta} \frac{\log(M)}{n} + r_1(n)^2 + r_1(n) \tilde{r}(n) + \frac{\log(M)}{\sqrt{n}} r_2(n) \right) \end{aligned}$$

Thus this upper bound, provides selection rates of the form:

$$\mathbb{E}[d(f_{\hat{k}}, f_0)] - \mathbb{E}[d(f_{\bar{k}}, f_0)] \lesssim 2\delta\mathbb{E}[d(f_{\bar{k}}, f_0)] + C(M_1, M_2) \left(\frac{1}{\delta} \frac{\log(M)}{n} + r_1(n)^2 + r_1(n) \tilde{r}(n) + \frac{\log(M)}{\sqrt{n}} r_2(n) \right)$$

As a side note, if we assume that the true nuisance parameter is known, i.e. g_0 is known, then $r_1(n) = r_2(n) = 0$, and we get the same type of results as in Van der Laan et al. (2007).

Invoking the theorem with the square loss that we used in our paper, we get:

$$\begin{aligned} d(f, f_0) &:= \|f - f_0\|^2 \\ r_1(n)^2 &= \Theta \left(\frac{\mathbb{E}[\text{BIAS}(X; \hat{g}) (f_{\hat{k}}(X) - f_0(X))]^2}{\mathbb{E}[\|f_{\hat{k}} - f_0\|^2]} \right) \stackrel{(i)}{\leq} \Theta \left(\frac{\mathbb{E}[\text{BIAS}(X; \hat{g})^2] \mathbb{E}[(f_{\hat{k}}(X) - f_0(X))^2]}{\mathbb{E}[\|f_{\hat{k}} - f_0\|^2]} \right) \\ &= \Theta \left(\frac{\mathbb{E}[\text{BIAS}(X; \hat{g})^2] \mathbb{E}[\|f_{\hat{k}} - f_0\|^2]}{\mathbb{E}[\|f_{\hat{k}} - f_0\|^2]} \right) = \Theta(\mathbb{E}[\text{BIAS}(X; \hat{g})^2]) \\ r_2(n)^2 &= \mathbb{E} \left[\max_j \left\{ \mathbb{E} \left[\mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2 | X] (f_j(X) - f_0(X))^2 \right] \right\} \right] = \Theta \left(\mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2] \right) \\ f_* &:= f_{\hat{k}}, \quad \tau_0 := f_0, \quad \hat{f} := f_{\hat{k}} \\ \tilde{r}(n) &= \|f_* - \tau_0\|^2 \end{aligned}$$

Moreover, inequality (i) is tight whenever $\text{BIAS}(X; \hat{g})$ is independent of $f_{\hat{k}}(X) - f_0(X)$, which cannot be excluded without further hard-to-justify assumptions. Therefore, the result of Van Der Laan and Dudoit (2003) yields bounds of the form:

$$\begin{aligned} \mathcal{S}(n) &:= \|\hat{f} - \tau_0\|^2 - \|f_* - \tau_0\|^2 \\ &\leq \Theta \left(\delta \|f_* - \tau_0\|^2 + \frac{1}{\delta} \frac{\log(M)}{n} + \mathbb{E}[\text{BIAS}(X; \hat{g})^2] + \sqrt{\mathbb{E}[\text{BIAS}(X; \hat{g})^2]} \|f_* - \tau_0\| + \frac{\log(M)}{\sqrt{n}} \mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2]^{\frac{1}{2}} \right) \end{aligned}$$

for any constant $\delta > 0$. Note that the optimal choice of δ , yields a rate of:

$$\begin{aligned} \mathcal{S}(n) &\leq \\ &\Theta \left(\|f_* - \tau_0\| \sqrt{\frac{\log(M)}{n}} + \mathbb{E}[\text{BIAS}(X; \hat{g})^2] + \sqrt{\mathbb{E}[\text{BIAS}(X; \hat{g})^2]} \|f_* - \tau_0\| + \frac{\log(M)}{\sqrt{n}} \mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2]^{\frac{1}{2}} \right) \end{aligned}$$

On the contrary, our Theorem 1 yields an oracle inequality of the form:

$$\mathcal{S}(n) \leq O \left(\epsilon + \frac{\log(M)}{n} + \mathbb{E}[\text{BIAS}(X; \hat{g})^2] \right) \quad \epsilon = \|f_* - \tau_0\|^2 - \|\bar{f}_* - \tau_0\|^2$$

We identify the following critical differences in the two results:

1. The term $\frac{\log(M)}{\sqrt{n}} \mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2]^{\frac{1}{2}}$ that appears in the result of Van Der Laan and Dudoit (2003) does not appear in our bound. This is crucial for the doubly robust application. Note that $\mathbb{E}[\text{BIAS}(X; \hat{g})^2]$ takes the form of the product of the errors of the propensity and the regression estimates, i.e. $\mathbb{E}[(\hat{\alpha} - \alpha_0)^2 (\hat{g} - g_0)^2]$. On the contrary the quantity $\mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2]^{\frac{1}{2}}$ does not take this product form. Thus, we would incur a dependence on the error of propensity and the regression, i.e. $\mathbb{E}[(\hat{\alpha} - \alpha_0)^2]^{\frac{1}{2}} + \mathbb{E}[(\hat{g} - g_0)^2]^{\frac{1}{2}}$. The latter wont converge faster than $1/\sqrt{n}$, even for parametric functions and thereby the estimation error of the regression and the propensity cannot be ignored in the selection error.

2. For any reasonable value of δ , the result of Van Der Laan and Dudoit (2003) would have a linear dependence on $\|f_* - \tau_0\|^2$. On the contrary our error only depends on the strictly smaller quantity $\epsilon = \|f_* - \tau_0\|^2 - \|\bar{f}_* - \tau_0\|^2$, where \bar{f}_* is the minizer of the square loss within the convex hull of the original set of functions. The quantity ϵ has the property that when there is a “dominant” model, such that the best in the convex hull puts almost all weight on one model, then ϵ is almost zero, while $\|f_* - \tau_0\|^2$ can be very far away from zero. Note that when $\|f_* - \tau_0\|$ is far from zero, then the optimal choice of δ , yields only a slow rate of $\sqrt{\log(M)/n}$ and not a fast rate of $\log(M)/n$.

Moreover, our main Q-aggregation Theorem 3 yields an oracle inequality of the form:

$$\mathcal{S}(n) \leq O\left(\frac{\log(M)}{n} + \mathbb{E}[\text{BIAS}(X; \hat{g})^2]\right)$$

We identify the following critical differences in the two results:

1. The term $\frac{\log(M)}{\sqrt{n}} \mathbb{E}\left[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2\right]^{\frac{1}{2}}$ does not appear in our bound, which is crucial for our doubly robust Q-aggregation results for the same reason as outlined in the bullet points above.
2. For any reasonable value of δ , the result of Van Der Laan and Dudoit (2003) would have a linear dependence on $\|f_* - \tau_0\|^2$. On the contrary our error bound incurs no dependence on the performance of the best model! Even if the best model is very far from the true function f_0 , we still compete with the best model at a $\log(M)/n + \mathbb{E}[\text{BIAS}(X; \hat{g})^2]$ rate.

B PROOFS FOR ERM-BASED SELECTION

B.1 Preliminaries

We will be invoking a generalization of orthogonal statistical learning theorem (Foster and Syrgkanis, 2023).

First, we describe the assumptions for the theorem. Let F and F' be two function spaces such that $F \subseteq F'$. Consider $f_* \in F$ and $f' \in F'$. Let $d(g, g_0)$ denote a distance metric in the function space G . The the remaining of the section, we will assume that the loss $\ell_{f,g}$ is C_b -lipschitz in $f(X)$.

ASSUMPTION 1 (Orthogonal Loss). *The population risk is Neyman Orthogonal, i.e. for all $\bar{f} \in \text{Star}(F, f')$ and $g \in G$:*

$$D_f D_g P \ell_{f', g_0}(Z)[g - g_0, f - f'] = 0$$

ASSUMPTION 2 (First Order Optimality). *The estimator f' for the population risk satisfies the first-order optimality condition:*

$$D_f P \ell_{f', g_0}(Z)[f - f'] \geq 0 \quad \forall f \in \text{Star}(F, f')$$

ASSUMPTION 3 (Higher Order Smoothness). *There exist constants β_1 and β_2 such that the following holds:*

- a) *Second order smoothness with respect to the target f . For all $f \in F$ and all $\bar{f} \in \text{Star}(F, f')$:*

$$D_{\bar{f}}^2 P \ell_{\bar{f}, g_0}(Z)[f - f', f - f'] \leq \beta_1 \|f - f'\|^2$$

- b) *Higher order smoothness. There exist $r \in [0, 1)$ such that for all $f \in \text{Star}(F, f')$, $g \in G$, and $\bar{g} \in \text{Star}(G, g_0)$:*

$$|D_{\bar{g}}^2 D_f P \ell_{f', \bar{g}}(Z)[f - f', g - g_0, g - g_0]| \leq \beta_2 \|f - f'\|^{1-r} d(g, g_0)^2$$

ASSUMPTION 4 (Strong Convexity). *The population risk is strongly convex with respect to the target function: there exist constants $\lambda, \kappa > 0$ such that for all $f \in F$ and an estimate $\hat{g} \in G$.*

$$D_{\bar{f}}^2 P \ell_{\bar{f}, \hat{g}}(Z)[f - f', f - f'] \geq \lambda \|f - f'\|^2 - \kappa d(\hat{g}, g_0)^{\frac{4}{1+r}} \quad \forall \bar{f} \in \text{Star}(F, f')$$

Alternatively, Assumption 1 and Assumption 3(b) may be replaced by the following assumption.

ASSUMPTION 5. *There exist $r \in [0, 1)$ and constant β_3 such that for all $f, \bar{f} \in \text{Star}(F, f')$, $g \in G$:*

$$|D_f P\ell(\bar{f}, g)[f - f'] - D_f P\ell(\bar{f}, g_0)[f - f']| \leq \beta_3 \|f - f'\|^{1-r} d(g, g_0)^2$$

Theorem 10 (Fast Rates Under Strong Convexity (Foster and Syrgkanis, 2023)). *Suppose Assumptions 1 to 4, or alternatively, Assumptions 2, 4, and 5 and Assumption 3(a), are satisfied for some $f' \in F'$ and function space $F \subseteq F'$. Then any estimator $\hat{f} \in F$ satisfies:*

$$P\ell_{\hat{f}, g_0}(Z) - P\ell_{f', g_0}(Z) \leq O\left(P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f', \hat{g}}(Z) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Moreover, if the estimator \hat{f} satisfies with probability at least $1 - \delta$,

$$P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f', \hat{g}}(Z) \leq \epsilon_n(\delta) \|\hat{f} - f'\| + \alpha_n(\delta)$$

for functions $\epsilon_n(\delta)$ and $\alpha_n(\delta)$, then we have, with probability at least $1 - \delta$

$$P\ell_{\hat{f}, g_0}(Z) - P\ell_{f', g_0}(Z) \leq O\left(\epsilon_n(\delta)^2 + \alpha_n(\delta) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Despite the small generalization made to consider any $f' \in F'$ that satisfies the first order optimality condition, the proof follows the original proof closely, and is included in Section J for completeness. We also showed that the same results can be obtained if Assumption 5 is satisfied instead of Assumptions 1 and 3.

Lemma 11. *Assume that the Assumptions 1 and 3(b) hold, or alternatively Assumption 5 holds, for some $f' \in F'$. Then for any $f_* \in F$ we have:*

$$P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f', \hat{g}}(Z) \leq P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f_*, \hat{g}}(Z) + 2(P\ell_{f_*, g_0}(Z) - P\ell_{f', g_0}(Z)) + O\left(\left((2/\lambda)^{\frac{2}{1+r}} + \kappa\right) d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Moreover, if ℓ is a square loss, we have:

$$P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f', \hat{g}}(Z) \leq P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f_*, \hat{g}}(Z) + 2(\|f_* - \tau_0\|^2 - \|f' - \tau_0\|^2) + O\left(d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Proof.

$$P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f', \hat{g}}(Z) = (P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f_*, \hat{g}}(Z)) + (P\ell_{f_*, \hat{g}}(Z) - P\ell_{f', \hat{g}}(Z))$$

Let us first consider the case where Assumptions 1 and 3(b) holds. For any g , taking a first order Taylor expansion yields:

$$P\ell_{f_*, g}(Z) - P\ell_{f', g}(Z) = D_f P\ell_{\bar{f}, g}(Z)[f_* - f']$$

for some $\bar{f} \in \text{Star}(F, f')$. Let us first consider the excess risk for evaluated at g_0 .

Taking a second order Taylor expansion for the $(P\ell_{f_*, \hat{g}}(Z) - P\ell_{f', \hat{g}}(Z))$ term at g_0 , we get:

$$\begin{aligned} (P\ell_{f_*, \hat{g}}(Z) - P\ell_{f', \hat{g}}(Z)) &= (P\ell_{f_*, g_0}(Z) - P\ell_{f', g_0}(Z)) + D_g D_f P\ell_{\bar{f}, g_0}(Z)[f_* - f', \hat{g} - g_0] \\ &\quad + \frac{1}{2} D_g^2 D_f P\ell_{\bar{f}, g_0}(Z)[f_* - f', \hat{g} - g_0, \hat{g} - g_0] \\ &= (P\ell_{f_*, g_0}(Z) - P\ell_{f', g_0}(Z)) + \frac{1}{2} D_g^2 D_f P\ell_{\bar{f}, g_0}(Z)[f_* - f', \hat{g} - g_0, \hat{g} - g_0] \\ &\hspace{15em} \text{(By Assumption 1)} \\ &\leq (P\ell_{f_*, g_0}(Z) - P\ell_{f', g_0}(Z)) + \frac{1}{2} D_g^2 D_f P\ell_{\bar{f}, g_0}(Z)[f_* - f', \hat{g} - g_0, \hat{g} - g_0] \\ &\leq (P\ell_{f_*, g_0}(Z) - P\ell_{f', g_0}(Z)) + \frac{\beta_2}{2} \|f_* - f'\|^{1-r} d(\hat{g}, g_0)^2 \quad \text{(By Assumption 3(b))} \\ &\leq (P\ell_{f_*, g_0}(Z) - P\ell_{f', g_0}(Z)) + \frac{\lambda}{2} \|f_* - f'\|^2 + O\left(\left((2/\lambda)^{\frac{2}{1+r}} d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)\right) \\ &\hspace{15em} \text{(Young's inequality)} \end{aligned}$$

Alternatively, if Assumption 5 holds, then it follows immediately from a first order Taylor expansion that

$$\begin{aligned} (Pl_{\hat{f}_*, \hat{g}}(Z) - Pl_{f', \hat{g}}(Z)) - (Pl_{f_*, g_0}(Z) - Pl_{f', g_0}(Z)) &= D_f Pl(\bar{f}, g)[\hat{f} - f'] - D_f Pl(\bar{f}, g_0)[\hat{f} - f'] \\ &\leq \beta_3 \|f - f'\|^{1-r} d(\hat{g}, g_0)^2 \\ &\leq \frac{\lambda}{2} \|f_* - f'\|^2 + O\left((2/\lambda)^{\frac{2}{1+r}} d(\hat{g}, g_0)^{\frac{4}{1+r}}\right) \end{aligned}$$

and the same result follows.

By the first order optimality of f' and strong convexity, we have that:

$$Pl_{f_*, g_0}(Z) - Pl_{f', g_0}(Z) \geq D_f Pl_{f', g_0}[f_* - f'] + \frac{\lambda}{2} \|f_* - f'\|^2 - \frac{\kappa}{2} d(\hat{g}, g_0)^{\frac{4}{1+r}} \geq \frac{\lambda}{2} \|f_* - f'\|^2 - \frac{\kappa}{2} d(\hat{g}, g_0)^{\frac{4}{1+r}}$$

Thus we deduce that:

$$Pl_{f_*, \hat{g}}(Z) - Pl_{f', \hat{g}}(Z) \leq 2(Pl_{f_*, g_0}(Z) - Pl_{f', g_0}(Z)) + O\left(\left((2/\lambda)^{\frac{2}{1+r}} + \kappa\right) d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

For square losses, $\lambda = 2$, $\kappa = 0$ and as we will show in Section F (by the realizability of τ_0 , i.e. $\mathbb{E}[Y(g_0) | X] = \tau_0(X)$):

$$(Pl_{f_*, g_0}(Z) - Pl_{f', g_0}(Z)) = \|f_* - \tau_0\|^2 - \|f' - \tau_0\|^2$$

So we get:

$$Pl_{\hat{f}, \hat{g}}(Z) - Pl_{f', \hat{g}}(Z) \leq (Pl_{\hat{f}, \hat{g}}(Z) - Pl_{f_*, \hat{g}}(Z)) + 2(\|f_* - \tau_0\|^2 - \|f' - \tau_0\|^2) + O\left(d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

□

Corollary 12 (Generalization of Theorem 10). *Suppose Assumptions 1 to 4, or alternatively, Assumptions 2, 4, and 5 and Assumption 3(a), are satisfied for some $f' \in F'$ and function space $F \subseteq F'$ and for some constants $\lambda, \beta_1, \beta_2, \beta_3, \kappa, r$. Let f_* be any function in F . Then any estimator $\hat{f} \in F$ satisfies*

$$Pl_{\hat{f}, g_0}(Z) - Pl_{f', g_0}(Z) \leq O\left(Pl_{\hat{f}, \hat{g}}(Z) - Pl_{f_*, \hat{g}}(Z) + (Pl_{f_*, g_0}(Z) - Pl_{f', g_0}(Z)) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

and for square losses:

$$Pl_{\hat{f}, g_0}(Z) - Pl_{f', g_0}(Z) \leq O\left(Pl_{\hat{f}, \hat{g}}(Z) - Pl_{f_*, \hat{g}}(Z) + \|f_* - \tau_0\|^2 - \|f' - \tau_0\|^2 + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Moreover, if the estimator \hat{f} satisfies w.p. $1 - \delta$

$$Pl_{\hat{f}, \hat{g}}(Z) - Pl_{f_*, \hat{g}}(Z) \leq \epsilon_n(\delta) \|\hat{f} - f_*\| + \alpha_n(\delta)$$

for functions $\epsilon_n(\delta)$ and $\alpha_n(\delta)$, then we have, with probability at least $1 - \delta$

$$Pl_{\hat{f}, g_0}(Z) - Pl_{f', g_0}(Z) \leq O\left((Pl_{f_*, g_0}(Z) - Pl_{f', g_0}(Z)) + \epsilon_n(\delta)^2 + \alpha_n(\delta) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

and for square losses:

$$Pl_{\hat{f}, g_0}(Z) - Pl_{f', g_0}(Z) \leq O\left(\|f_* - \tau_0\|^2 - \|f' - \tau_0\|^2 + \epsilon_n(\delta)^2 + \alpha_n(\delta) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Proof. The first result follows from a direct application of Lemma 11.

Next,

$$\begin{aligned} Pl_{\hat{f}, \hat{g}}(Z) - Pl_{f_*, \hat{g}}(Z) &\leq \epsilon_n(\delta) \|\hat{f} - f_*\| + \alpha_n(\delta) \\ &\leq \epsilon_n(\delta) (\|\hat{f} - f'\| + \|f_* - f'\|) + \alpha_n(\delta) \\ &\leq \epsilon_n(\delta) \|\hat{f} - f'\| + \frac{1}{2} \epsilon_n(\delta)^2 + \frac{1}{2} \|f_* - f'\|^2 + \alpha_n(\delta) \end{aligned}$$

By the first order optimality of f' and strong convexity, we have that:

$$Pl_{f_*,g_0}(Z) - Pl_{f',g_0}(Z) \geq D_f Pl_{f',g_0}[f_* - f'] + \lambda \|f_* - f'\|^2 - \kappa d(g, g_0)^{\frac{4}{1+r}} \geq \lambda \|f_* - f'\|^2 - \kappa d(\hat{g}, g_0)^{\frac{4}{1+r}}$$

Thus:

$$Pl_{\hat{f},\hat{g}}(Z) - Pl_{f_*,\hat{g}}(Z) \leq \epsilon_n(\delta) \|\hat{f} - f'\| + \frac{1}{2} \epsilon_n(\delta)^2 + \alpha_n(\delta) + O\left(Pl_{f_*,g_0}(Z) - Pl_{f',g_0}(Z) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Applying Lemma 11 yields:

$$\begin{aligned} Pl_{\hat{f},\hat{g}}(Z) - Pl_{f',\hat{g}}(Z) &\leq (Pl_{\hat{f},\hat{g}}(Z) - Pl_{f_*,\hat{g}}(Z)) + O\left((Pl_{f_*,g_0}(Z) - Pl_{f',g_0}(Z)) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right) \\ &\leq \epsilon_n(\delta) \|\hat{f} - f'\| + O\left(\epsilon_n(\delta)^2 + \alpha_n(\delta) + (Pl_{f_*,g_0}(Z) - Pl_{f',g_0}(Z)) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right) \end{aligned}$$

Finally, invoking Theorem 10, we get:

$$Pl_{\hat{f},g_0}(Z) - Pl_{f',g_0}(Z) \leq O\left(\epsilon_n(\delta)^2 + \alpha_n(\delta) + (Pl_{f_*,g_0}(Z) - Pl_{f',g_0}(Z)) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Similarly for square losses. □

When strong convexity is not satisfied, slower oracle rates can be shown.

ASSUMPTION 6 (Universal Orthogonality). *For all $\bar{f} \in \text{Star}(F, f_*) + \text{Star}(F - f_*, 0)$,*

$$D_g D_f Pl_{\bar{f},g_0}(Z)[f - f_*, g - g_0] = 0 \quad \forall g \in G, f \in F$$

ASSUMPTION 7. *The derivatives $D_g^2 Pl_{f,g}(Z)$ and $D_f^2 D_g Pl_{f,g}(Z)$ are continuous. Furthermore, there exists a constant β such that for all $f \in \text{Star}(F, f_*)$ and $\bar{g} \in \text{Star}(G, g_0)$,*

$$|D_g^2 Pl_{f,\bar{g}}[g - g_0, g - g_0]| \leq \beta d(g, g_0)^2 \quad \forall g \in G$$

Similarly, Assumptions 6 to 7 can be replaced by the following assumption:

ASSUMPTION 8. *There exist a constant β_4 such that for all $f, \bar{f} \in \text{Star}(F, f')$, $g \in G$:*

$$|D_f Pl(\bar{f}, g)[f - f'] - D_f Pl(\bar{f}, g_0)[f - f']| \leq \beta_4 d(g, g_0)^2$$

Theorem 13 (Slow Rates without Strong Convexity). *Suppose that f_* satisfies Assumptions 6 to 7, or alternatively, Assumption 8 holds, then with probability at least $1 - \delta$, the target estimator \hat{f} has the following excess risk bound:*

$$Pl_{\hat{f},g_0}(Z) - Pl_{f',g_0}(Z) \leq O\left(Pl_{\hat{f},\hat{g}}(Z) - Pl_{f',\hat{g}}(Z) + d(g, g_0)^2\right)$$

B.2 Best-ERM: Proof of Theorem 1 of the Main Paper

First, we verify that Assumptions 2, 4, and 5 are satisfied for the square losses, so that Corollary 12 can be invoked.

Consider f' to be the minimizer of the population risk in the function space $F' = \text{co}(F)$. Since we can rewrite $\ell_{f,g_0}(Z) = \mathbb{E}[(f - \tau_0)^2] + \text{Var}(\hat{Y}(g_0))$, it is straight forward to see that Assumption 2 is satisfied.

Moreover, for any function g , we have:

$$D_f^2 Pl_{\bar{f},g_0}(Z)[f - f', f - f'] = \mathbb{E}[2(f - f')^2] = 2\|f - f'\|^2$$

Therefore, Assumption 4 are satisfied with $r = 0$ and $\kappa = 0$.

In Section F, we showed that:

$$|(Pl_{f,g}(Z) - Pl_{f',g}(Z)) - (Pl_{f,g_0}(Z) - Pl_{f',g_0}(Z))| = 2\sqrt{\mathbb{E}[\text{BIAS}(\hat{g})^2]} \|f_{\theta^*} - f_{\hat{\theta}}\|$$

Thus, Assumption 5 is satisfied with $r = 0$, $\beta_3 = 2$, and $d(\hat{g}, g_0) = \mathbb{E} [\text{BIAS}(\hat{g})^2]^{\frac{1}{4}}$.

Next, we need to bound $P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f_*, \hat{g}}(Z)$.

Invoking Lemma 14 from Foster and Syrgkanis (2023), with probability at least $1 - \delta$, we can bound the centered excess loss for ERM over any function class F by:

$$(P - P_n)(\ell_{\hat{f}, \hat{g}}(Z) - \ell_{f_*, \hat{g}}(Z)) \leq O\left(\delta_n \|\hat{f} - f_*\| + \delta_n^2\right)$$

where for any $f_* \in F$, $\delta_n = r_n + \sqrt{\frac{\log(1/\delta)}{n}}$, and r_n is the critical radius of $\text{Star}(F - f_*, 0)$. In particular, we consider f_* to be the minimizer of $\|f_* - f'\|$ in F . By definition of \hat{f} , $P_n(\ell_{\hat{f}, \hat{g}}(Z) - \ell_{f_*, \hat{g}}(Z)) < 0$, so we have:

$$P(\ell_{\hat{f}, \hat{g}}(Z) - \ell_{f_*, \hat{g}}(Z)) \leq (P - P_n)(\ell_{\hat{f}, \hat{g}}(Z) - \ell_{f_*, \hat{g}}(Z)) \leq O\left(\delta_n \|\hat{f} - f_*\| + \delta_n^2\right)$$

As a result, we can instantiate Corollary 12 with $r = 0$, and $d(\hat{g}, g_0) = \mathbb{E} [\text{BIAS}(\hat{g})^2]^{\frac{1}{4}}$. Moreover, we can let $\epsilon_n = O(\delta_n)$ and $\alpha_n = O(\delta_n^2)$, to obtain:

$$\begin{aligned} \|\hat{f} - \tau_0\|^2 - \|f' - \tau_0\|^2 &= P\ell_{\hat{f}, g_0}(Z) - P\ell_{f', g_0}(Z) \\ &\leq O\left(\delta_n^2 + \|f_* - \tau_0\|^2 - \|f' - \tau_0\|^2 + \mathbb{E} [\text{BIAS}(\hat{g})^2]\right) \end{aligned}$$

Applying the fact that, for $F = \{f_1, f_2, \dots, f_M\}$, the critical radius of the $\text{Star}(F - f_*, 0)$ is $r_n^2 = \Theta\left(\frac{\log(\max\{M, n\})}{n}\right)$, we get:

$$\|\hat{f} - \tau_0\|^2 - \|f' - \tau_0\|^2 \leq O\left(\frac{\log(\max\{M, n\}/\delta)}{n} + \|f_* - \tau_0\|^2 - \|f' - \tau_0\|^2 + \mathbb{E} [\text{BIAS}(\hat{g})^2]\right)$$

Since by definition, f' is equivalent to \bar{f}_* , we obtain a stronger result than in Theorem 1 in the main text, which bounds the regret with respect to the optimal convex combination of models instead of the best model:

$$\begin{aligned} \mathcal{S}(\hat{\theta}) &= \|f_{\hat{\theta}} - \tau_0\|^2 - \|f_{\theta_*} - \tau_0\|^2 \\ &\leq \|f_{\hat{\theta}} - \tau_0\|^2 - \|\bar{f}_{\theta_*} - \tau_0\|^2 \\ &\leq O\left(\frac{\log(\max\{M, n\}/\delta)}{n} + \|f_* - \tau_0\|^2 - \|f' - \tau_0\|^2 + \mathbb{E} [\text{BIAS}(\hat{g})^2]\right) \end{aligned}$$

B.3 Convex-ERM: Proof of Theorem 2 of the Main Paper

For ERM over the convex hull, we can follow the same arguments to invoke Theorem 10 with $r = 0$, $\beta_3 = 2$, and $d(\hat{g}, g_0) = \mathbb{E} [\text{BIAS}(\hat{g})^2]^{\frac{1}{4}}$ as in the previous section, with $f_* = f' = \bar{f}_*$ and $F' = F = \text{co}(\{f_1, f_2, \dots, f_M\})$, we get that:

$$\|\hat{f} - \tau_0\|^2 - \|f_* - \tau_0\|^2 \leq O\left(\delta_n^2 + \mathbb{E} [\text{BIAS}(\hat{g})^2]\right)$$

By metric entropy arguments, the critical radius for this $\text{Star}(F - f_*, 0)$ is $r_n = \Theta\left(\sqrt{\frac{M \log(n)}{n}}\right)$. Combining, we get:

$$\|\hat{f} - \tau_0\|^2 - \|f_* - \tau_0\|^2 \leq O\left(\frac{M \log(n/\delta)}{n} + \mathbb{E} [\text{BIAS}(\hat{g})^2]\right)$$

On the other hand, a different nuisance error term can be obtained by invoking Theorem 13. First, we show that square losses satisfy Assumption 8. Assuming f is uniformly bounded by U for all $f \in F$, we have that:

$$\begin{aligned} |D_f P\ell(\bar{f}, g)[f - f'] - D_f P\ell(\bar{f}, g_0)[f - f']| &= |\mathbb{E}[(\bar{f} - \hat{Y}(g))(f - f') - (\bar{f} - \hat{Y}(g_0))(f - f')]| \\ &= |\mathbb{E}[(\hat{Y}(g) - \hat{Y}(g_0))(f - f')]| \\ &= 2|\mathbb{E}[\text{BIAS}(\hat{g})(\hat{f} - f')]| \\ &\leq 4U\mathbb{E}[|\text{BIAS}(\hat{g})|] \end{aligned}$$

Thus, Assumption 8 is satisfied with $\beta_4 = 4U$ and $d(\hat{g}, g_0) = \mathbb{E}[\|\text{BIAS}(\hat{g})\|]^{\frac{1}{2}}$.

Next, we need a generalization bound that does not involve $\|\hat{f} - f'\|$ terms. With probability at least $1 - \delta$:

$$\begin{aligned} (P - P_n)(\ell_{\hat{f}, \hat{g}} - \ell_{f', \hat{g}}) &\leq C_b(P - P_n)(\hat{f} - f') && \text{(Since we assume } \ell \text{ is } C_b \text{ Lipschitz in } f(X)) \\ &\leq O\left(R(F) + \sqrt{\frac{\log(1/\delta)}{n}}\right) && \text{(By standard generalization arguments)} \end{aligned}$$

where $R(F)$ is the Rademacher complexity of the function space F .

Thus invoking Theorem 13, we get:

$$\begin{aligned} \|\hat{f} - \tau_0\|^2 - \|f_* - \tau_0\|^2 &\leq O\left(R(F) + \sqrt{\frac{\log(1/\delta)}{n}} + \mathbb{E}[\|\text{BIAS}(\hat{g})\|]\right) \\ &\leq O\left(\sqrt{\frac{\log(M)}{n}} + \sqrt{\frac{\log(1/\delta)}{n}} + \mathbb{E}[\|\text{BIAS}(\hat{g})\|]\right) \\ &\quad \text{(Lemma 26.11 in Shalev-Shwartz and Ben-David (2014))} \\ &\leq O\left(\sqrt{\frac{\log(M/\delta)}{n}} + \mathbb{E}[\|\text{BIAS}(\hat{g})\|]\right) \end{aligned}$$

As with the Best-ERM case, this bounds the regret with respect to the optimal convex combination of models instead of the best model, which is a stronger result than in Theorem 2 in the main text.

C MIXED BIAS PROPERTIES

C.1 Doubly Robust Targets for CATE under Conditional Exogeneity

$$\begin{aligned} \mathbb{E}[\hat{Y}(g_0) - \hat{Y}(\hat{g}) \mid X] &= \mathbb{E}[a_0(D, W) h_0(D, W) \mid X] - \mathbb{E}[a_0(D, W) \hat{h}(D, W) \mid X] \\ &\quad + \mathbb{E}[\hat{a}(D, W) (h_0(D, W) - \hat{h}(D, W)) \mid X] \\ &= \mathbb{E}[(a_0(D, W) - \hat{a}(D, W)) (h_0(D, W) - \hat{h}(D, W)) \mid X] \\ &= \mathbb{E}[\hat{q}(D, W) \mid X] \end{aligned}$$

C.2 Doubly Robust Targets for CATE with Instruments

$$\begin{aligned} \mathbb{E}[\hat{Y}(\hat{g}) - \hat{Y}(g_0) \mid X] &= \hat{\tau}(X) + \frac{\mathbb{E}[Y\tilde{Z} \mid X] - \hat{\tau}(X)\mathbb{E}[D\tilde{Z} \mid X]}{\hat{\beta}(X)} - \tau_0(X) \\ &= \hat{\tau}(X) + \frac{\tau_0(X)\beta_0(X) - \hat{\tau}(X)\beta_0(X)}{\hat{\beta}(X)} - \tau_0(X) \\ &= \hat{\tau}(X) - \tau_0(X) + (\tau_0(X) - \hat{\tau}(X))\frac{\beta_0(X)}{\hat{\beta}(X)} \\ &= (\tau_0(X) - \hat{\tau}(X))\left(\frac{\beta_0(X)}{\hat{\beta}(X)} - 1\right) \\ &= (\tau_0(X) - \hat{\tau}(X))\frac{\beta_0(X) - \hat{\beta}(X)}{\hat{\beta}(X)} \end{aligned}$$

D EXPERIMENT DETAILS

D.1 Candidate Models

We consider 8 meta-learners to be the candidate models for this study. The estimators $\hat{\mu}(X, D)$, $\hat{\mu}_0(X)$, $\hat{\mu}_1(X)$, $\hat{\pi}$ are fitted using XGboost regressors and classifiers.

T-learner

$$\hat{\tau}_T(X) = \hat{\mu}_1(X) - \hat{\mu}_0(X)$$

S-learner

$$\hat{\tau}_S(X) = \hat{\mu}(X, D = 1) - \hat{\mu}(X, D = 0)$$

Inverse Propensity Weighted Learner

$$\hat{\tau}_{IPW} = \arg \min_{f \in F} \sum_i \left(\frac{D_i Y_i}{\hat{\pi}(X_i)} - \frac{(1 - D_i) Y_i}{1 - \hat{\pi}(X_i)} - f(X_i) \right)^2$$

X-learner (Künzel et al., 2019) Using the potential outcome estimators, we first compute the imputed treatment effects:

$$\begin{aligned} \hat{\tau}^0 &= \arg \min_{f \in F} \sum_{i, D_i=0} (\hat{\mu}_1(X_i) - Y_i - f(X_i))^2 \\ \hat{\tau}^1 &= \arg \min_{f \in F} \sum_{i, D_i=1} (Y_i - \hat{\mu}_0(X_i) - f(X_i))^2 \end{aligned}$$

The final estimator is an weighted average of the two imputed treatment effects:

$$\hat{\tau}_{IPW}(X) = \hat{\pi}(X) \hat{\tau}^0 + (1 - \hat{\pi}(X)) \hat{\tau}^1$$

DR-learner (Foster and Syrgkanis, 2023; Kennedy, 2020) The doubly robust potential outcome is formulated by the addition of the inverse propensity weighted residual term:

$$Y_d^{DR} = \hat{\mu}(X, d) + \frac{Y - \hat{\mu}(X, D)}{\hat{\pi}_d(X)} \cdot \mathbb{1}(D = d)$$

The final estimator is:

$$\hat{\tau}_{DR} := \arg \min_{f \in F} \sum_i (Y_{i,1}^{DR} - Y_{i,0}^{DR} - f(X_i))^2$$

R-learner (Nie and Wager, 2021) The R-learner employs an additional nuisance estimator, $\hat{m}(X)$, for the conditional response surface $\mathbb{E}[Y | X]$. Let $\tilde{Y} = Y - \hat{m}(X)$ and $\tilde{D} = D - \hat{\pi}(X)$ denote the residuals. Then, the final estimator is:

$$\hat{\tau}_R := \arg \min_{f \in F} \sum_i \left(\tilde{Y}_i - f(X_i) \cdot \tilde{D}_i \right)^2$$

Doubly Robust X-learner The DRX-learner follows the doubly robust framework, but replaces $\hat{\mu}(X, d)$ by their X-learner analogues:

$$\begin{aligned} \hat{g}_0(X) &= \hat{\mu}(X, 0)(1 - \hat{\pi}(X)) + (\hat{\mu}(X, 1) - \hat{\tau}^0) \hat{\pi}(X) \\ \hat{g}_1(X) &= \hat{\mu}(X, 1)(1 - \hat{\pi}(X)) + (\hat{\mu}(X, 0) - \hat{\tau}^1) \hat{\pi}(X) \end{aligned}$$

DAX-learner This is a combination of the X-learner strategy and the Domain-Adaptation learner proposed in the EconML library Battocchi et al. (2019) that we introduce in this paper. In the context of the X-Learner, we train a model $\hat{\tau}^1$ on the treated data and then we use it to calculate $\hat{\tau}(X) = \hat{\tau}^1(X) (1 - \hat{\pi}(X)) + \hat{\tau}^0(X) \hat{\pi}(X)$ on

all the data points. Thus in this case, when measuring the quality of the downstream CATE estimate $\hat{\tau}$ in the final step of the X-Learner, we care about the quality of $\hat{\tau}^1$ as measured by the metric:

$$\mathbb{E}_X [(1 - \hat{\pi}(X))^2 (\hat{\tau}^1(X) - \tau_0(X))^2]$$

On the contrary, the regression estimator for $\hat{\tau}^1$ would be guaranteeing:

$$\mathbb{E}_{X|D=1} [(\hat{\tau}^1(X) - \tau_0(X))^2]$$

We can correct this covariate shift if instead we minimized the weighted square loss:

$$\mathbb{E}_{X|D=1} \left[\frac{1}{\hat{\pi}(X)} (1 - \hat{\pi}(X))^2 (\hat{\tau}^1(X) - \tau_0(X))^2 \right]$$

which corresponds to a density ratio weighted square loss. We can achieve this by using a sample-weighted regression oracle with weights $(1 - \hat{\pi}(X))^2 \frac{1}{\hat{\pi}(X)}$, when training $\hat{\tau}^1$. Similarly, we should use weights $\hat{\pi}(X)^2 \frac{1}{1 - \hat{\pi}(X)}$, when training $\hat{\tau}^0$. Then we proceed exactly as in the X-learner. We will call this the Domain-Adapted-X-Learner (DAX-Learner).

D.2 Data Generating Processes for Synthetic Datasets

The details of each DGP is summarized in Table 3.

Table 3: Summary of Data Generating Processes.

DGP	$\pi(X)$	$m(X)$	$\tau(X)$
1	0.05	$0.11(0.6 \leq X \leq 0.8)$	0.5
2	0.05	$0.11(0.6 \leq X \leq 0.8)$	$0.5 + 0.11(0.2 \leq X \leq 0.4)$
3	$0.5 - 0.491(0.3 \leq X \leq 0.6)$	$0.11(0.6 \leq X \leq 0.8)$	$0.5X^2$
4	$0.5 - 0.491(0.3 \leq X \leq 0.6)$	$0.11(0.6 \leq X \leq 0.8)$	$0.5X^2 \mathbf{1}(X \leq 0.6) + 0.18$
5	0.05	0.1	$0.5 \mathbf{1}(0.5 \leq X \leq 0.8)$
6	0.95	0.1	$0.5 \mathbf{1}(0.5 \leq X \leq 0.8)$

D.3 Semi-synthetic Datasets

The details dataset (after pre-processing) used for Semi-synthetic experiments are summarized in the Table 4.

Table 4: Summary of Semi-synthetic Datasets.

	401k	welfare	STAR	Poverty	Criteo
Number of Entries	9716	12907	5220	2480	13979592
Number of Covariates	14	42	15	24	12
Percentage Treated	37.18%	51.33%	30.23%	50.89%	85.00%

D.4 Additional Results

The RMSE regret for the candidate models as well as model selection, convex stacking, and Q-aggregation is summarized in tables 6 to 8. Each cell in the table reports [mean \pm st.dev.] median (95%) of RMSE Regret over the 100 experiments for each semi-synthetic dataset or DGP. The last two rows presents the percentage decrease in the mean and median RMSE regret for Q-aggregation with respect to convex stacking and model selection using the DR-loss.

Table 5: Mean RMSE over Fully Synthetic Datasets

	S	T	IPS	X	DR	R	DRX	DAX	Q-agg
DGP 1	0.026	0.032	0.0606	0.0104	0.0174	0.00362	0.0174	0.0179	0.0163
DGP 2	0.0137	0.0164	0.0454	0.00984	0.00973	0.00966	0.0098	0.00975	0.00883
DGP 3	0.037	0.0141	0.0318	0.0107	0.0131	0.0151	0.012	0.0179	0.0088
DGP 4	0.0166	0.0079	0.075	0.0208	0.0102	0.0145	0.0204	0.0156	0.00819
DGP 5	0.122	0.13	0.104	0.13	0.0519	0.135	0.0526	0.0306	0.0326
DGP 6	0.159	0.00479	0.058	0.175	0.0355	0.189	0.104	0.0899	0.00955




Figure 1: Median RMSE Regret for Each DGP.

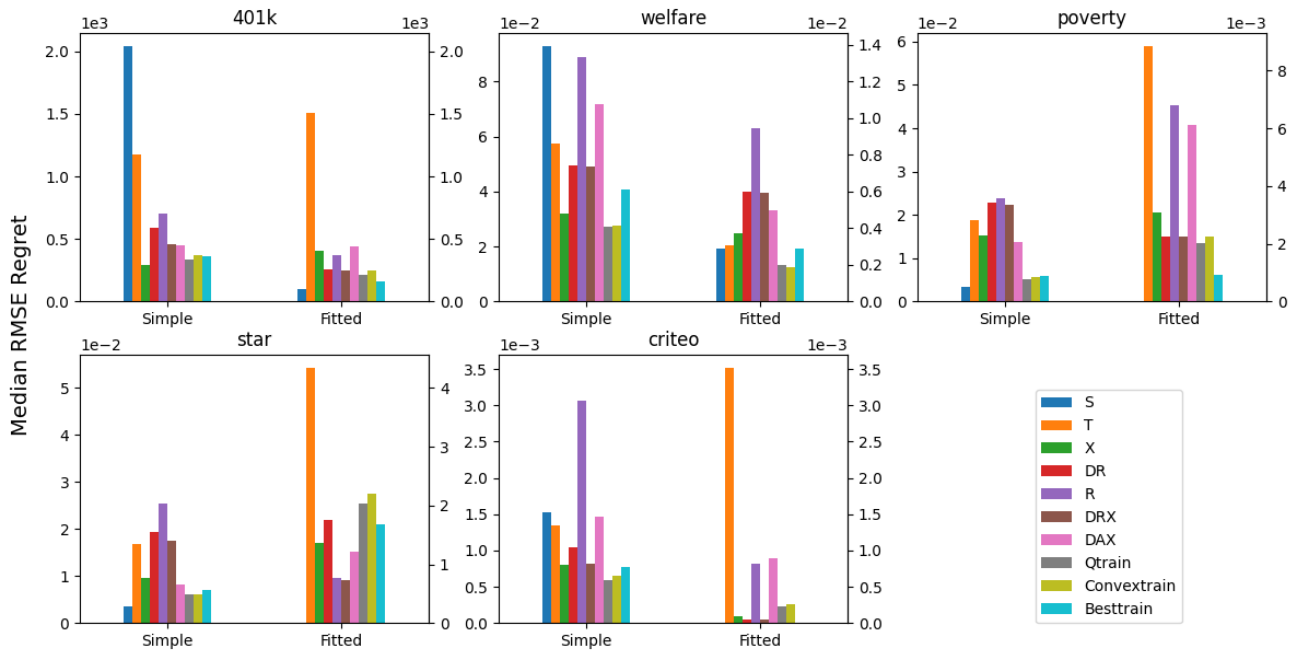


Figure 2: Median RMSE Regret for Each Semi-synthetic Dataset.

Table 6: RMSE Regret For Fully-synthetic Datasets.

	DGP 1	DGP 2	DGP 3	DGP 4	DGP 5	DGP 6
S	[0.026 ± 0.012] 0.023 (0.047)	[0.014 ± 0.009] 0.012 (0.032)	[0.037 ± 0.010] 0.039 (0.051)	[0.017 ± 0.006] 0.015 (0.027)	[0.122 ± 0.034] 0.115 (0.186)	[0.159 ± 0.026] 0.157 (0.194)
T	[0.032 ± 0.007] 0.032 (0.042)	[0.016 ± 0.007] 0.016 (0.029)	[0.014 ± 0.011] 0.012 (0.036)	[0.008 ± 0.006] 0.007 (0.020)	[0.130 ± 0.027] 0.130 (0.168)	[0.005 ± 0.006] 0.002 (0.018)
IPS	[0.061 ± 0.047] 0.052 (0.143)	[0.045 ± 0.038] 0.032 (0.118)	[0.032 ± 0.008] 0.032 (0.047)	[0.075 ± 0.020] 0.073 (0.112)	[0.104 ± 0.044] 0.111 (0.164)	[0.058 ± 0.021] 0.058 (0.094)
X	[0.010 ± 0.009] 0.009 (0.025)	[0.010 ± 0.006] 0.009 (0.021)	[0.011 ± 0.009] 0.009 (0.023)	[0.021 ± 0.008] 0.021 (0.033)	[0.130 ± 0.026] 0.130 (0.168)	[0.175 ± 0.024] 0.183 (0.204)
DR	[0.017 ± 0.011] 0.016 (0.037)	[0.010 ± 0.007] 0.008 (0.022)	[0.013 ± 0.011] 0.011 (0.035)	[0.010 ± 0.007] 0.010 (0.022)	[0.052 ± 0.034] 0.048 (0.111)	[0.036 ± 0.015] 0.034 (0.063)
R	[0.004 ± 0.005] 0.001 (0.016)	[0.010 ± 0.005] 0.009 (0.019)	[0.015 ± 0.012] 0.010 (0.037)	[0.015 ± 0.008] 0.015 (0.029)	[0.135 ± 0.026] 0.136 (0.174)	[0.189 ± 0.024] 0.200 (0.215)
DRX	[0.017 ± 0.012] 0.015 (0.039)	[0.010 ± 0.008] 0.007 (0.026)	[0.012 ± 0.009] 0.010 (0.036)	[0.020 ± 0.008] 0.020 (0.032)	[0.053 ± 0.034] 0.048 (0.111)	[0.104 ± 0.035] 0.095 (0.176)
DAX	[0.018 ± 0.014] 0.012 (0.045)	[0.010 ± 0.008] 0.008 (0.025)	[0.018 ± 0.016] 0.013 (0.047)	[0.016 ± 0.016] 0.012 (0.042)	[0.031 ± 0.027] 0.024 (0.076)	[0.090 ± 0.028] 0.083 (0.140)
Besttrain	[0.017 ± 0.013] 0.014 (0.040)	[0.011 ± 0.008] 0.010 (0.027)	[0.012 ± 0.011] 0.009 (0.036)	[0.010 ± 0.009] 0.008 (0.021)	[0.036 ± 0.031] 0.032 (0.089)	[0.008 ± 0.013] 0.003 (0.041)
Convextrain	[0.017 ± 0.011] 0.016 (0.039)	[0.009 ± 0.008] 0.007 (0.027)	[0.008 ± 0.009] 0.006 (0.025)	[0.008 ± 0.007] 0.007 (0.020)	[0.033 ± 0.025] 0.028 (0.083)	[0.011 ± 0.009] 0.009 (0.027)
Qtrain	[0.016 ± 0.011] 0.015 (0.038)	[0.009 ± 0.007] 0.007 (0.023)	[0.009 ± 0.009] 0.006 (0.025)	[0.008 ± 0.007] 0.007 (0.021)	[0.033 ± 0.025] 0.028 (0.083)	[0.010 ± 0.010] 0.006 (0.026)
% Decrease w.r.t Best	1.264 % -7.700 %	22.081 % 40.148 %	41.138 % 44.073 %	20.952 % 25.093 %	11.236 % 11.446 %	-13.003 % -49.579 %
% Decrease w.r.t Convex	2.324 % 2.740 %	0.119 % -6.044 %	-4.316 % -8.569 %	0.526 % 5.734 %	0.080 % -0.636 %	12.474 % 35.252 %

Table 7: RMSE Regret For Semi-synthetic Datasets with Simple CATE.

	401k	welfare	poverty	star	criteo
S	[2042 ± 411.7] 2042 (2677)	[0.0947 ± 0.0248] 0.0929 (0.1346)	[0.0050 ± 0.0050] 0.0034 (0.0139)	[0.0055 ± 0.0059] 0.0035 (0.0174)	[0.00151 ± 0.00044] 0.0015 (0.00223)
T	[1172 ± 229.5] 1172 (1572)	[0.0546 ± 0.0267] 0.0573 (0.1035)	[0.0193 ± 0.0072] 0.0188 (0.0320)	[0.0224 ± 0.0211] 0.0169 (0.0670)	[0.00136 ± 0.00054] 0.0013 (0.00231)
IPS	[5845 ± 1489] 5513 (8983)	[0.7937 ± 0.0893] 0.7757 (0.9694)	[0.1739 ± 0.0373] 0.1731 (0.2329)	[0.6294 ± 0.2361] 0.5938 (1.0945)	[0.05433 ± 0.02190] 0.0459 (0.09577)
X	[302.6 ± 155.6] 289.2 (582.4)	[0.0337 ± 0.0170] 0.0319 (0.0649)	[0.0159 ± 0.0064] 0.0152 (0.0267)	[0.0145 ± 0.0158] 0.0096 (0.0490)	[0.00079 ± 0.00032] 0.0008 (0.00135)
DR	[581 ± 225.8] 587.6 (994.3)	[0.0522 ± 0.0195] 0.0494 (0.0876)	[0.0241 ± 0.0068] 0.0229 (0.0375)	[0.0206 ± 0.0165] 0.0194 (0.0431)	[0.00105 ± 0.00055] 0.0010 (0.00193)
R	[727.5 ± 217.9] 704.6 (1096)	[0.0872 ± 0.0174] 0.0889 (0.1157)	[0.0248 ± 0.0072] 0.0239 (0.0378)	[0.0280 ± 0.0176] 0.0254 (0.0540)	[0.00302 ± 0.00057] 0.0031 (0.00390)
DRX	[460.5 ± 194] 460.3 (801.7)	[0.0510 ± 0.0188] 0.0489 (0.0832)	[0.0240 ± 0.0069] 0.0223 (0.0390)	[0.0198 ± 0.0157] 0.0176 (0.0411)	[0.00082 ± 0.00040] 0.0008 (0.00161)
DAX	[446.6 ± 179.7] 444.6 (741.2)	[0.0727 ± 0.0173] 0.0718 (0.1047)	[0.0147 ± 0.0068] 0.0138 (0.0260)	[0.0133 ± 0.0170] 0.0083 (0.0389)	[0.00145 ± 0.00030] 0.0015 (0.00190)
Besttrain	[355 ± 147.9] 362.4 (585.7)	[0.0390 ± 0.0181] 0.0406 (0.0722)	[0.0075 ± 0.0068] 0.0060 (0.0194)	[0.0107 ± 0.0096] 0.0070 (0.0287)	[0.00081 ± 0.00040] 0.0008 (0.00161)
Convextrain	[362.3 ± 157.1] 371.5 (634.4)	[0.0289 ± 0.0156] 0.0275 (0.0557)	[0.0066 ± 0.0049] 0.0057 (0.0166)	[0.0095 ± 0.0093] 0.0062 (0.0284)	[0.00066 ± 0.00032] 0.0006 (0.00122)
Qtrain	[341.1 ± 157.1] 331.3 (634.2)	[0.0281 ± 0.0156] 0.0272 (0.0537)	[0.0061 ± 0.0049] 0.0051 (0.0165)	[0.0095 ± 0.0093] 0.0061 (0.0284)	[0.00060 ± 0.00032] 0.0006 (0.00112)
% Decrease w.r.t Best	4.0692 % 9.4008 %	38.7569 % 49.2038 %	22.8005 % 17.7918 %	13.0432 % 15.7684 %	34.8879 % 30.0913 %
% Decrease w.r.t Convex	6.2148 % 12.1341 %	2.6238 % 1.0101 %	8.1161 % 13.0241 %	0.2770 % 2.3870 %	9.9629 % 8.5184 %

Table 8: RMSE Regret For Semi-synthetic Datasets with Fitted CATE.

	401k	welfare	poverty	star	criteo
S	[115.1 ± 104.1] 98.59 (306.3)	[0.0029 ± 0.0018] 0.0029 (0.0061)	[0.0002 ± 0.0004] 0.0000 (0.0013)	[0.0843 ± 0.1903] 0.0000 (0.4518)	[0.00000 ± 0.00001] 0.0000 (0.00001)
T	[1526 ± 340.2] 1511 (2052)	[0.0033 ± 0.0014] 0.0031 (0.0056)	[0.0088 ± 0.0013] 0.0088 (0.0110)	[4.3250 ± 0.6061] 4.3450 (5.1982)	[0.00334 ± 0.00089] 0.0035 (0.00478)
IPS	[1643 ± 1291] 1212 (4731)	[0.0343 ± 0.0059] 0.0343 (0.0437)	[0.0219 ± 0.0151] 0.0181 (0.0575)	[345.6718 ± 130.7644] 321.2090 (563.5089)	[0.00227 ± 0.00107] 0.0024 (0.00376)
X	[431.7 ± 271.9] 400.8 (919.6)	[0.0039 ± 0.0017] 0.0037 (0.0069)	[0.0032 ± 0.0010] 0.0031 (0.0048)	[1.3916 ± 0.4319] 1.3608 (2.1141)	[0.00013 ± 0.00011] 0.0001 (0.00035)
DR	[321.8 ± 284.4] 259.5 (882.3)	[0.0063 ± 0.0021] 0.0060 (0.0103)	[0.0025 ± 0.0014] 0.0022 (0.0048)	[1.6791 ± 0.8769] 1.7604 (3.0720)	[0.00006 ± 0.00004] 0.0000 (0.00015)
R	[410.2 ± 216.7] 374 (837.5)	[0.0095 ± 0.0020] 0.0094 (0.0129)	[0.0068 ± 0.0011] 0.0068 (0.0086)	[0.8925 ± 0.4380] 0.7595 (1.7322)	[0.00081 ± 0.00008] 0.0008 (0.00094)
DRX	[314.5 ± 279.5] 251.4 (842.7)	[0.0063 ± 0.0021] 0.0059 (0.0103)	[0.0025 ± 0.0014] 0.0022 (0.0048)	[0.8536 ± 0.6130] 0.7277 (2.0083)	[0.00006 ± 0.00004] 0.0001 (0.00013)
DAX	[501.6 ± 254.7] 437.1 (1063)	[0.0052 ± 0.0015] 0.0050 (0.0080)	[0.0062 ± 0.0011] 0.0061 (0.0082)	[1.2216 ± 0.4937] 1.2192 (2.0556)	[0.00090 ± 0.00021] 0.0009 (0.00123)
Besttrain	[257.5 ± 262.1] 159.9 (812.1)	[0.0029 ± 0.0015] 0.0029 (0.0053)	[0.0017 ± 0.0020] 0.0009 (0.0057)	[2.0018 ± 1.4700] 1.6779 (4.7187)	[0.00002 ± 0.00003] 0.0000 (0.00010)
Convextrain	[266.2 ± 180.9] 245.3 (577.4)	[0.0019 ± 0.0012] 0.0019 (0.0040)	[0.0023 ± 0.0014] 0.0022 (0.0048)	[2.1629 ± 0.8906] 2.2012 (3.5904)	[0.00027 ± 0.00007] 0.0003 (0.00038)
Qtrain	[245.6 ± 183] 215.8 (579.5)	[0.0019 ± 0.0012] 0.0020 (0.0041)	[0.0021 ± 0.0014] 0.0020 (0.0047)	[2.0705 ± 0.9286] 2.0370 (3.5692)	[0.00024 ± 0.00006] 0.0002 (0.00035)
% Decrease w.r.t Best	4.8579 % -25.8868 %	50.4544 % 46.1057 %	-17.3363 % -55.4313 %	-3.3152 % -17.6290 %	-91.0409 % -100.0000 %
% Decrease w.r.t Convex	8.4122 % 13.6626 %	-0.7623 % -4.4831 %	8.7545 % 9.9212 %	4.4656 % 8.0570 %	10.7980 % 11.0042 %

E PROOF OF MAIN THEOREM

In this section, we will be proving Theorem 3 in the main paper, which we also reproduce here:

Theorem 14 (Main Theorem). *Assume that for some function $\text{ERROR}(\hat{g})$:*

$$P(\tilde{\ell}_{\theta, g_0}(Z) - \tilde{\ell}_{\theta', g_0}(Z)) - P(\tilde{\ell}_{\theta, \hat{g}}(Z) - \tilde{\ell}_{\theta', \hat{g}}(Z)) \leq \text{ERROR}(\hat{g}) \|f_{\theta} - f_{\theta'}\|^{\gamma}$$

for some $\gamma \in [0, 1]$ and $\ell(f, g_0)$ is σ -strongly convex in expectation with respect to f . Moreover, assume that the loss $\ell(Z; f(X), \hat{g})$ is $C_b(\hat{g})$ -Lipschitz with respect to $f(X)$, and both the loss and the candidate functions are uniformly bounded in $[-b, b]$ almost surely. For any $\mu \leq \min\{\nu, 1 - \nu\} \frac{\sigma}{14}$ and $\beta \geq \left\{ \frac{8C_b^2(1-\nu)^2}{\mu}, 8\sqrt{3}bC_b(1-\nu), \frac{2C_b\nu(\nu C_b + \mu b)}{\mu} \right\}$, w.p. $1 - \delta$:

$$R(\hat{\theta}, g_0) \leq \min_{j=1}^M \left(R(e_j, g_0) + \frac{\beta}{n} \log(1/\pi_j) \right) + O \left(\frac{\log(1/\delta) \max\{\frac{1}{\mu} C_b(\hat{g})^2, b C_b(\hat{g})\}}{n} \right) + \left(\frac{1}{\mu} \right)^{\frac{\gamma}{2-\gamma}} \text{ERROR}(\hat{g})^{\frac{2}{2-\gamma}}$$

We remind some key definitions and assumptions from the main text. We assume that the function ℓ is σ -strongly convex in f in expectation, i.e.:

$$P(\ell(Z; f'(X), g_0) - \ell(Z; f(X), g_0)) \geq P(\langle \nabla_f \ell(Z; f(X), g_0), f'(X) - f(X) \rangle) + \frac{\sigma}{2} \|f - f'\|^2$$

We consider the modified loss, which penalizes the weight put on each model, based on the individual model performance:

$$\tilde{\ell}_{\theta, g}(Z) = (1 - \nu)\ell_{\theta, g}(Z) + \nu \sum_j \theta_j \ell_{e_j, g}(Z) \quad \tilde{R}(\theta, g) = P\tilde{\ell}_{\theta, g}(Z)$$

We also define with:

$$K(\theta) := \sum_j \theta_j \log(1/\pi_j)$$

Then the Q-aggregation ensemble is defined as:

$$\hat{\theta} = \arg \min_{\theta} P_n \tilde{\ell}_{\theta, \hat{g}}(Z) + \frac{\beta}{n} K(\theta)$$

We introduce some population quantities that will play a key role in our analysis. Define:

$$V(\theta) := \sum_j \theta_j \|f_j - f_\theta\|^2$$

$$H(\theta, g) := \tilde{R}(\theta, g) + \mu V(\theta) + \frac{\beta}{n} K(\theta)$$

and define:

$$\theta^* = \arg \min_{\theta} \tilde{R}(\theta, g_0) + \mu V(\theta) + \frac{\beta}{n} K(\theta) = \arg \min_{\theta} H(\theta, g_0)$$

Then we introduce the shorthand notation and invoke the error decomposition assumption:

$$\mathcal{E}(\hat{g}) := P(\tilde{\ell}_{\hat{\theta}, g_0}(Z) - \tilde{\ell}_{\theta^*, g_0}(Z)) - P(\tilde{\ell}_{\hat{\theta}, \hat{g}}(Z) - \tilde{\ell}_{\theta^*, \hat{g}}(Z)) \leq \text{ERROR}(\hat{g}) \|f_{\hat{\theta}} - f_{\theta^*}\|^\gamma$$

Lemma 15. *If $\ell(Z; f(X), g)$ is $\sigma(g)$ -strongly convex in expectation with respect to f , then for any $\theta, \theta' \in \Theta$:*

$$R(\theta, g) - R(\theta', g) \geq \langle \nabla R(\theta', g), \theta - \theta' \rangle + \frac{\sigma(g)}{2} \|f_\theta - f_{\theta'}\|^2$$

Proof. Let F be the random column vector defined as $F := [f_1(X) \cdots f_M(X)]$. Note that by the chain rule:

$$\begin{aligned} \langle \nabla_{\theta} \ell(Z; f'_{\theta}(X), g), \theta - \theta' \rangle &= \langle (\nabla_{\theta} (f'_{\theta}(X)))^{\top} \nabla_f \ell(Z; f_{\theta'}(X), g), \theta - \theta' \rangle \\ &= \langle F \nabla_f \ell(Z; f_{\theta'}(X), g), \theta - \theta' \rangle \\ &= \nabla_f \ell(Z; f_{\theta'}(X), g)^{\top} F^{\top} (\theta - \theta') \\ &= \langle \nabla_f \ell(Z; f_{\theta'}(X), g), F^{\top} (\theta - \theta') \rangle \\ &= \langle \nabla_f \ell(Z; f_{\theta'}(X), g), f_{\theta}(X) - f'_{\theta'}(X) \rangle \end{aligned}$$

Taking expectation over Z , we have:

$$\begin{aligned} \langle \nabla_{\theta} R(\theta'), \theta - \theta' \rangle &= \langle \nabla_{\theta} \mathbb{E}[\ell(Z; f'_{\theta'}(X), g)], \theta - \theta' \rangle \\ &= \mathbb{E}[\langle \nabla_{\theta} \ell(Z; f'_{\theta'}(X), g), \theta - \theta' \rangle] \\ &= \mathbb{E}[\langle \nabla_f \ell(Z; f_{\theta'}(X), g), f_{\theta}(X) - f'_{\theta'}(X) \rangle] \end{aligned}$$

The result then follows by the definition of $\sigma(g)$ -strong convexity in expectation of the function $\ell(f, g)$ and the definition of $R(\theta, g)$. \square

Lemma 16. *If $R(\theta, g)$ is $\sigma(g)$ -strongly convex in expectation, then for any $\theta \in \Theta$:*

$$\tilde{R}(\theta, g) \geq R(\theta, g) + \frac{\nu \sigma(g)}{2} V(\theta)$$

Proof. By σ -strong convexity of the loss and Lemma 15, we have that:

$$R(e_j, g) - R(\theta, g) \geq \langle \nabla R(\theta, g), e_j - \theta \rangle + \frac{\sigma(g)}{2} \|f_j - f_\theta\|^2$$

Multiplying each equation by θ_j and summing across j , yields:

$$R(\theta, g) \leq \sum_j \theta_j R(e_j, g) - \frac{\sigma(g)}{2} V(\theta)$$

where we used the fact that $\sum_j \theta_j = 1$ and $\sum_j \theta_j \langle \nabla R(\theta, g), e_j - \theta \rangle = 0$.

Thus we have that:

$$\begin{aligned}\tilde{R}(\theta, g) &= (1 - \nu)R(\theta, g) + \nu \sum_j \theta_j R(e_j, g) \\ &\geq (1 - \nu)R(\theta, g) + \nu R(\theta, g) + \frac{\nu\sigma(g)}{2}V(\theta) \\ &= R(\theta, g) + \frac{\nu\sigma(g)}{2}V(\theta)\end{aligned}$$

□

Lemma 17. *If $\ell(Z; f(X), g_0)$ is σ -strongly convex in f , in expectation, and $\theta_* = \arg \min_{\theta \in \Theta} H(\theta, g_0)$, then for any $\theta \in \Theta$:*

$$H(\theta, g_0) - H(\theta^*, g_0) \geq \left(\frac{(1 - \nu)\sigma}{2} - \mu \right) \|f_\theta - f_{\theta^*}\|^2$$

Proof. By Lemma 15 we have:

$$R(\theta, g_0) - R(\theta', g_0) \geq \langle \nabla R(\theta', g_0), \theta - \theta' \rangle + \frac{\sigma}{2} \|f_\theta - f_{\theta'}\|^2 \quad (3)$$

Now let $q(Z; \theta, g) = \sum_{j=1}^M \theta_j \ell(Z; f_{e_j}(X), g)$ and $L_e(g) = [\ell(Z; f_{e_1}(X), g) \cdots \ell(Z; f_{e_M}(X), g)]^T$ and let $Q(\theta, g) = \mathbb{E}[q(Z; \theta, g)]$. We have that:

$$q(Z; \theta, g) - q(Z; \theta', g) = \sum_{j=1}^M \ell(Z; f_{e_j}(X), g)(\theta_j - \theta'_j) = \langle L_e(g), \theta - \theta' \rangle = \langle \nabla q(Z; \theta', g), \theta - \theta' \rangle$$

By taking the expectation over Z :

$$Q(\theta, g) - Q(\theta', g) = \sum_{j=1}^M \theta_j R(e_j, g) - \sum_{j=1}^M \theta'_j R(e_j, g) = \langle \nabla Q(\theta', g), \theta - \theta' \rangle \quad (4)$$

By adding Equations (3) and (4) we have that:

$$\begin{aligned}\tilde{R}(\theta, g_0) - \tilde{R}(\theta', g_0) &= (1 - \nu)(R(\theta, g_0) - R(\theta', g_0)) + \nu(Q(\theta, g_0) - Q(\theta', g_0)) \\ &\geq (1 - \nu)\langle \nabla R(\theta', g_0), \theta - \theta' \rangle + \nu\langle \nabla Q(\theta', g_0), \theta - \theta' \rangle + \frac{(1 - \nu)\sigma}{2} \|f_\theta - f_{\theta'}\|^2 \\ &\geq \langle \nabla \tilde{R}(\theta', g_0), \theta - \theta' \rangle + \frac{(1 - \nu)\sigma}{2} \|f_\theta - f_{\theta'}\|^2\end{aligned}$$

meaning that \tilde{R} will also be $(1 - \nu)\sigma$ -strongly convex in θ . We also have that for any θ, θ' :²

$$V(\theta) - V(\theta') = \langle \nabla V(\theta'), \theta - \theta' \rangle - \|f_\theta - f_{\theta'}\|_2^2$$

which implies that $V(\theta)$ is at most 2-strongly concave. We also have:

$$K(\theta) - K(\theta') = \sum_j (\theta_j - \theta'_j) \log(1/\pi_j) = \langle \nabla K(\theta'), \theta - \theta' \rangle$$

Adding the three inequalities and substituting θ' with θ^* we have:

$$H(\theta, g_0) - H(\theta^*, g_0) \geq \langle \nabla H(\theta^*, g_0), \theta - \theta^* \rangle + \left(\frac{(1 - \nu)\sigma}{2} - \mu \right) \|f_\theta - f_{\theta^*}\|^2$$

Note that by the optimality of θ^* over the convex set Θ we have that $\langle \nabla H(\theta^*, g_0), \theta - \theta^* \rangle \geq 0$, which immediately yields the statement we want to prove. □

²Since the function $V(\theta)$ can be interpreted as the expected variance of $f_j(X)$ over the distribution of j . Thus with simple manipulations it can be shown to be equal to $V(\theta) = \sum_j \theta_j \|f_j\|^2 - \|f_\theta\|^2$. Since this is a quadratic in θ function, with Hessian $\mathbb{E}[FF^T]$, where F is the vector $(f_1(X), \dots, f_j(X))$, the above equality follows by an exact second order Taylor expansion and the observation that $(\theta - \theta')^T \mathbb{E}[FF^T](\theta - \theta') = \|f_\theta - f_{\theta'}\|^2$

Lemma 18. *Define the centered and offset empirical processes:*

$$\begin{aligned} Z_{n,1} &= (1 - \nu)(P - P_n)(\ell_{\hat{\theta}, \hat{g}}(Z) - \ell_{\theta^*, \hat{g}}(Z)) - \mu \sum_j \hat{\theta}_j \|f_j - f_{\theta^*}\|^2 - \frac{1}{s} K(\hat{\theta}) \\ Z_{n,2} &= \nu(P - P_n) \sum_j (\hat{\theta}_j - \theta_j^*) \ell_{e_j, \hat{g}}(Z) - \mu \hat{\theta}^\top H \theta^* - \frac{1}{s} K(\hat{\theta}) \end{aligned} \quad H_{ij} = \|f_i - f_j\|^2$$

where s is a positive number. Assume that:

$$\mathcal{E}(\hat{g}) := P(\tilde{\ell}_{\hat{\theta}, g_0}(Z) - \tilde{\ell}_{\theta^*, g_0}(Z)) - P(\tilde{\ell}_{\hat{\theta}, \hat{g}}(Z) - \tilde{\ell}_{\theta^*, \hat{g}}(Z)) \leq \text{ERROR}(\hat{g}) \|f_\theta - f_{\theta^*}\|$$

and $\ell(Z; f(X), g_0)$ is σ -strongly convex in expectation with respect to f . For $\beta \geq \frac{4n}{s}$ and $\mu \leq \min\{\nu, 1 - \nu\} \frac{\sigma}{14}$:

$$R(\hat{\theta}) \leq \min_j [R(e_j, g_0) + \frac{\beta}{n} \log(1/\pi_j)] + 2(Z_{n,1} + Z_{n,2}) + \frac{1}{\mu} \text{ERROR}(\hat{g})^2$$

Proof.

$$\begin{aligned} \tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0) &= P(\tilde{\ell}_{\hat{\theta}, g_0}(Z) - \tilde{\ell}_{\theta^*, g_0}(Z)) \\ &= P(\tilde{\ell}_{\hat{\theta}, \hat{g}}(Z) - \tilde{\ell}_{\theta^*, \hat{g}}(Z)) + \mathcal{E}(\hat{g}) \\ &\leq (P - P_n)(\tilde{\ell}_{\hat{\theta}, \hat{g}}(Z) - \tilde{\ell}_{\theta^*, \hat{g}}(Z)) + \mathcal{E}(\hat{g}) - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) \\ &\quad \text{(By optimality of } \hat{\theta}, P_n(\tilde{\ell}_{\hat{\theta}, \hat{g}}(Z) - \tilde{\ell}_{\theta^*, \hat{g}}(Z)) + \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) \leq 0) \\ &= \mu \sum_j \hat{\theta}_j \|f_j - f_{\theta^*}\|^2 + \mu \hat{\theta}^\top H \theta^* + Z_{n,1} + Z_{n,2} + \mathcal{E}(\hat{g}) \\ &\quad - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + \frac{2}{s} K(\hat{\theta}) \end{aligned} \quad \text{(definition of } Z_{n,1}, Z_{n,2})$$

By simple algebraic manipulations note that for any $\theta \in \Theta$:

$$\begin{aligned} \sum_j \hat{\theta}_j \|f_j - f_{\theta^*}\|^2 &= V(\hat{\theta}) + \|f_{\hat{\theta}} - f_{\theta^*}\|^2 \\ \hat{\theta}^\top H \theta^* &= V(\hat{\theta}) + V(\theta^*) + \|f_{\hat{\theta}} - f_{\theta^*}\|^2 \end{aligned}$$

Letting $Z_n = Z_{n,1} + Z_{n,2}$, we have:

$$\begin{aligned} \tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0) &\leq 2\mu V(\hat{\theta}) + \mu V(\theta^*) + 2\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 + Z_n + \mathcal{E}(\hat{g}) - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + \frac{2}{s} K(\hat{\theta}) \\ &\leq 2\mu V(\hat{\theta}) + \mu V(\theta^*) + 2\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 + Z_n + \text{ERROR}(\hat{g}) \|f_{\hat{\theta}} - f_{\theta^*}\|^\gamma \\ &\quad - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + \frac{2}{s} K(\hat{\theta}) \end{aligned}$$

By Young's inequality, for any $a \geq 0$ and $b \geq 0$, we have $ab \leq \frac{a^p}{p} + \frac{b^q}{q}$, such that $\frac{1}{p} + \frac{1}{q} = 1$. Taking $a = \left(\frac{\gamma}{2\mu}\right)^{\gamma/2} \text{ERROR}(\hat{g})$ and $b = \left(\frac{2\mu}{\gamma}\right)^{\gamma/2} \|f_{\hat{\theta}} - f_{\theta^*}\|^\gamma$ and $q = 2/\gamma$ and $p = \frac{1}{1-1/q} = \frac{1}{1-\gamma/2} = \frac{2}{2-\gamma}$

$$\text{ERROR}(\hat{g}) \|f_{\hat{\theta}} - f_{\theta^*}\|^\gamma \leq \underbrace{\frac{2}{2-\gamma} \left(\frac{\gamma}{2\mu}\right)^{\frac{\gamma}{2-\gamma}} \text{ERROR}(\hat{g})^{\frac{2}{2-\gamma}} + \mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2}_{\mathcal{A}(\hat{g})}$$

Thus we can derive:

$$\begin{aligned} \tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0) &\leq 2\mu V(\hat{\theta}) + \mu V(\theta^*) + 3\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 + Z_n + \mathcal{A}(\hat{g}) \\ &\quad - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + \frac{2}{s} K(\hat{\theta}) \end{aligned} \quad (5)$$

By Lemma 17, we have:

$$\begin{aligned} \tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0) &\geq \mu(V(\theta^*) - V(\hat{\theta})) + \frac{\beta}{n}(K(\theta^*) - K(\hat{\theta})) + \left(\frac{(1-\nu)\sigma}{2} - \mu\right) \|f_{\hat{\theta}} - f_{\theta^*}\|^2 \\ &\geq \mu(V(\theta^*) - V(\hat{\theta})) + \frac{\beta}{n}(K(\theta^*) - K(\hat{\theta})) + 6\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 \\ &\quad \text{(Since: } \mu \leq \frac{(1-\nu)\sigma}{14} \implies \frac{(1-\nu)\sigma}{2} \geq 7\mu) \end{aligned}$$

Combining the latter upper and lower bound on $\tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0)$, we have:

$$3\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 \leq 3\mu V(\hat{\theta}) + Z_n + \frac{2}{s}K(\hat{\theta}) + \mathcal{A}(\hat{g})$$

Plugging the bound on $3\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2$ back into Equation (5) we have:

$$\begin{aligned} \tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0) &\leq 5\mu V(\hat{\theta}) + \mu V(\theta^*) + 2Z_n + 2\mathcal{A}(\hat{g}) \\ &\quad + \frac{\beta}{n}(K(\theta^*) - K(\hat{\theta})) + \frac{4}{s}K(\hat{\theta}) \end{aligned}$$

By Lemma 16, we have:

$$\begin{aligned} R(\hat{\theta}, g_0) &\leq \tilde{R}(\theta^*, g_0) + \mu V(\theta^*) + \left(5\mu - \frac{\nu\sigma(g_0)}{2}\right) V(\hat{\theta}) + 2Z_n + 2\mathcal{A}(\hat{g}) \\ &\quad + \frac{\beta}{n}K(\theta^*) + \left(\frac{4}{s} - \frac{\beta}{n}\right) K(\hat{\theta}) \\ &\leq \tilde{R}(\theta^*, g_0) + \mu V(\theta^*) + \frac{\beta}{n}K(\theta^*) + 2Z_n + 2\mathcal{A}(\hat{g}) \quad \text{(since } \mu \leq \frac{\nu\sigma}{10}, V(\hat{\theta}) > 0, \beta \geq \frac{4n}{s} \text{ and } K(\hat{\theta}) \geq 0) \\ &= \min_{\theta \in \Theta} [\tilde{R}(\theta, g_0) + \mu V(\theta) + \frac{\beta}{n}K(\theta^*)] + 2Z_n + 2\mathcal{A}(\hat{g}) \quad \text{(Since } \theta_* \text{ minimizes } H(\theta, g_0)) \\ &\leq \min_j [\tilde{R}(e_j, g_0) + \mu V(e_j) + \frac{\beta}{n} \log(1/\pi_j)] + 2Z_n + 2\mathcal{A}(\hat{g}) \\ &= \min_j [R(e_j, g_0) + \frac{\beta}{n} \log(1/\pi_j)] + 2Z_n + 2\mathcal{A}(\hat{g}) \quad \text{(Since } \tilde{R}(e_j, g_0) = R(e_j, g_0) \text{ and } V(e_j) = 0 \text{ for all } e_j) \end{aligned}$$

□

E.1 Concentration of offset empirical processes

Lemma 19. *Assuming that the loss $\ell(f(X), \hat{g})$ is C_b lipschitz in $f(X)$ and all $f_i(X)$ are bounded in $[-b, b]$, almost surely, then for any positive $s \leq \min \left\{ \frac{\mu n}{(2C_b(1-\nu))^2}, \frac{n}{2\sqrt{3}bC_b(1-\nu)} \right\}$:*

$$\Pr \left(Z_{n,1} \geq \frac{\log(1/\delta)}{s} \right) \leq \delta$$

Proof. By Lemma 21 it suffices to show that:

$$I := \mathbb{E} [\exp \{sZ_{n,1}\}] \leq 1$$

By symmetrization (sym.) and contraction (contr.):

$$\begin{aligned}
 I &= \mathbb{E} \left[\exp \left\{ s \left((1-\nu)(P - P_n)(\ell_{\hat{\theta}}(Z) - \ell_{\theta^*}(Z)) - \mu \sum_j \hat{\theta}_j \|f_j - f_{\theta^*}\|^2 - \frac{1}{s} \sum_j \hat{\theta}_j \log(1/\pi_j) \right) \right\} \right] \\
 &\leq \mathbb{E} \left[\exp \left\{ s \max_{\theta} \left((1-\nu)(P - P_n)(\ell_{\theta}(Z) - \ell_{\theta^*}(Z)) - \mu \sum_j \theta_j \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \sum_j \theta_j \log(\pi_j) \right) \right\} \right] \\
 &\leq \mathbb{E} \left[\exp \left\{ s \max_{\theta} \left(2(1-\nu)P_{n,\epsilon}(\ell_{\theta}(Z) - \ell_{\theta^*}(Z)) - \mu \sum_j \theta_j \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \sum_j \theta_j \log(\pi_j) \right) \right\} \right] \quad (\text{sym.}) \\
 &\leq \mathbb{E} \left[\exp \left\{ s \max_{\theta} \left(2(1-\nu)C_b P_{n,\epsilon}(f_{\theta}(Z) - f_{\theta^*}(Z)) - \mu \sum_j \theta_j \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \sum_j \theta_j \log(\pi_j) \right) \right\} \right] \quad (\text{contr.})
 \end{aligned}$$

Since the quantity:

$$2(1-\nu)C_b P_{n,\epsilon}(f_{\theta}(Z) - f_{\theta^*}(Z)) - \mu \sum_j \theta_j \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \sum_j \theta_j \log(\pi_j)$$

is a linear function of θ (recall that $f_{\theta} = \sum_j \theta_j f_j$), the maximum over θ is attained at one of e_1, \dots, e_M . Thus:

$$\begin{aligned}
 I &\leq \mathbb{E} \left[\exp \left\{ s \max_j \left(2(1-\nu)C_b P_{n,\epsilon}(f_j(Z) - f_{\theta^*}(Z)) - \mu \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \log(\pi_j) \right) \right\} \right] \\
 &\leq \mathbb{E} \left[\max_j \pi_j \exp \left\{ s \left(2(1-\nu)C_b P_{n,\epsilon}(f_j(Z) - f_{\theta^*}(Z)) - \mu \|f_j - f_{\theta^*}\|^2 \right) \right\} \right] \\
 &\leq \mathbb{E} \left[\sum_j \pi_j \exp \left\{ s \left(2(1-\nu)C_b P_{n,\epsilon}(f_j(Z) - f_{\theta^*}(Z)) - \mu \|f_j - f_{\theta^*}\|^2 \right) \right\} \right] \\
 &= \sum_j \pi_j \mathbb{E} \left[\exp \left\{ s \left(2(1-\nu)C_b P_{n,\epsilon}(f_j(Z) - f_{\theta^*}(Z)) - \mu \|f_j - f_{\theta^*}\|^2 \right) \right\} \right] \\
 &= \sum_j \pi_j \mathbb{E}_{Z_{1:n}} \mathbb{E}_{\epsilon} \left[\exp \left\{ s \left(2(1-\nu)C_b P_{n,\epsilon}(f_j(Z) - f_{\theta^*}(Z)) - \mu \|f_j - f_{\theta^*}\|^2 \right) \right\} \right] \\
 &= \sum_j \pi_j \mathbb{E}_{Z_{1:n}} \left[\mathbb{E}_{\epsilon} \left[\exp \left\{ s 2(1-\nu)C_b P_{n,\epsilon}(f_j(Z) - f_{\theta^*}(Z)) \right\} \right] \exp \left\{ -s\mu \|f_j - f_{\theta^*}\|^2 \right\} \right]
 \end{aligned}$$

By Hoeffding's inequality in Lemma 22, we can bound the first term in the product:

$$\mathbb{E}_{\epsilon} \left[\exp \left\{ s 2(1-\nu)C_b P_{n,\epsilon}(f_j(Z) - f_{\theta^*}(Z)) \right\} \right] \leq \exp \left\{ \frac{2s^2}{n} (1-\nu)^2 C_b^2 P_n(f_j(Z) - f_{\theta^*}(Z))^2 \right\}$$

Thus we have:

$$\begin{aligned}
 I &\leq \sum_j \pi_j \mathbb{E}_{Z_{1:n}} \left[\exp \left\{ \frac{2s^2}{n} (1-\nu)^2 C_b^2 P_n(f_j(Z) - f_{\theta^*}(Z))^2 \right\} \exp \left\{ -s\mu \|f_j - f_{\theta^*}\|^2 \right\} \right] \\
 &= \sum_j \pi_j \mathbb{E}_{Z_{1:n}} \left[\exp \left\{ \frac{2s^2}{n} (1-\nu)^2 C_b^2 P_n(f_j(Z) - f_{\theta^*}(Z))^2 \right\} \exp \left\{ -s\mu P(f_j(Z) - f_{\theta^*}(Z))^2 \right\} \right] \\
 &= \sum_j \pi_j \mathbb{E}_{Z_{1:n}} \left[\exp \left\{ \frac{2s^2(1-\nu)^2 C_b^2}{n} \left(P_n - \frac{\mu n}{2s(1-\nu)^2 C_b^2} P \right) (f_j(Z) - f_{\theta^*}(Z))^2 \right\} \right]
 \end{aligned}$$

Since s is small enough such that: $\frac{\mu n}{2s(1-\nu)^2 C_b^2} \geq 2$ (which is satisfied by our assumptions on s), we have that:

$$\begin{aligned} I &\leq \sum_j \pi_j \mathbb{E}_{Z_{1:n}} \left[\exp \left\{ \frac{2s^2(1-\nu)^2 C_b^2}{n} (P_n - 2P) (f_j(Z) - f_{\theta^*}(Z))^2 \right\} \right] \\ &= \sum_j \pi_j \mathbb{E}_{Z_{1:n}} \left[\exp \left\{ \frac{2s^2(1-\nu)^2 C_b^2}{n} \left((P_n - P) (f_j(Z) - f_{\theta^*}(Z))^2 - P(f_j(Z) - f_{\theta^*}(Z))^2 \right) \right\} \right] \end{aligned}$$

Since $f_j(Z) \in [-b, b]$, we have that $f_j(Z) - f_{\theta^*}(Z) \in [-2b, 2b]$. Thus:

$$(f_j(Z) - f_{\theta^*}(Z))^2 \geq \frac{(f_j(Z) - f_{\theta^*}(Z))^2}{4b^2} (f_j(Z) - f_{\theta^*}(Z))^2 = \frac{1}{4b^2} (f_j(Z) - f_{\theta^*}(Z))^4.$$

Thus:

$$I \leq \sum_j \pi_j \mathbb{E}_{Z_{1:n}} \left[\exp \left\{ \frac{2s^2(1-\nu)^2 C_b^2}{n} \left((P_n - P) (f_j(Z) - f_{\theta^*}(Z))^2 - \frac{1}{4b^2} P(f_j(Z) - f_{\theta^*}(Z))^4 \right) \right\} \right]$$

It suffices to show that each of the summands is upper bounded by 1. Then their weighted average is also upper bounded by 1.

Applying Lemma 23 with:

$$\begin{aligned} V_i &= (f_j(Z_i) - f_{\theta^*}(Z_i))^2 \\ c &= 4b^2 \\ \lambda &= \frac{2s^2(1-\nu)^2 C_b^2}{n^2} \\ c_0 &= \frac{1}{4b^2} \end{aligned}$$

we have that as long as:

$$\lambda := \frac{2s^2(1-\nu)^2 C_b^2}{n^2} \leq \frac{2c_0}{1+2c_0c} =: \frac{2 \cdot 1/4b^2}{1+2} = \frac{1}{6b^2}$$

Then:

$$\mathbb{E}_{Z_{1:n}} \left[\exp \left\{ \frac{2s^2(1-\nu)^2 C_b^2}{n} \left((P_n - P) (f_j(Z) - f_{\theta^*}(Z))^2 - \frac{1}{4b^2} P(f_j(Z) - f_{\theta^*}(Z))^4 \right) \right\} \right] \leq 1$$

Thus we need that:

$$s \leq \frac{n}{2\sqrt{3}b(1-\nu)C_b}$$

which is satisfied by our assumptions on s . □

Lemma 20. *Assuming that the loss $\ell(f(X), \hat{g})$ is C_b lipschitz in $f(X)$ and bounded in $[-b, b]$, for any positive $s \leq \frac{2\mu n}{C_b \nu (C_b + 4\mu b)}$:*

$$\Pr \left(Z_{n,2} \geq \frac{\log(1/\delta)}{s} \right) \leq \delta$$

Proof. It suffices to show that:

$$I := \mathbb{E} [\exp \{sZ_{n,2}\}] \leq 1$$

By definition of $Z_{n,2}$:

$$\begin{aligned} I &= \mathbb{E} \left[\exp \left\{ s \left(\nu(P - P_n) \sum_j (\hat{\theta}_j - \theta_j^*) \ell_{e_j}(Z) - \mu \hat{\theta}^\top H \theta^* - \frac{1}{s} \sum_j \hat{\theta}_j \log(1/\pi_j) \right) \right\} \right] \\ &= \mathbb{E} \left[\exp \left\{ s \left(\nu(P - P_n) \sum_j \hat{\theta}_j \ell_{e_j}(Z) + \frac{1}{s} \sum_j \hat{\theta}_j \log(\pi_j) \right) \right\} \exp \left\{ s \left(-\nu(P - P_n) \sum_j \theta_j^* \ell_{e_j}(Z) - \mu \hat{\theta}^\top H \theta^* \right) \right\} \right] \end{aligned}$$

Note that:

$$\begin{aligned} -\nu(P - P_n) \sum_j \theta_j^* \ell_{e_j}(Z) - \mu \hat{\theta}^\top H \theta^* &= \sum_j \theta_j^* \left(-\nu(P - P_n) \ell_{e_j}(Z) - \mu \hat{\theta}^\top H_{\cdot,j} \right) \\ &= \sum_j \theta_j^* \underbrace{\left(-\nu(P - P_n) \ell_{e_j}(Z) - \mu \sum_k \hat{\theta}_k \|f_k - f_j\|^2 \right)}_{\Phi_j} \end{aligned}$$

Thus:

$$I = \mathbb{E} \left[\exp \left\{ s \left(\nu(P - P_n) \sum_j \hat{\theta}_j \ell_{e_j}(Z) + \frac{1}{s} \sum_j \hat{\theta}_j \log(\pi_j) \right) \right\} \exp \left\{ \sum_j \theta_j^* s \Phi_j \right\} \right]$$

By Jensen's inequality:

$$\exp \left\{ \sum_j \theta_j^* s \Phi_j \right\} \leq \sum_j \theta_j^* \exp \{ s \Phi_j \}$$

Thus:

$$\begin{aligned} I &\leq \sum_j \theta_j^* \mathbb{E} \left[\exp \left\{ s \left(\nu(P - P_n) \sum_k \hat{\theta}_k \ell_{e_k}(Z) + \frac{1}{s} \sum_k \hat{\theta}_k \log(\pi_k) \right) \right\} \exp \{ s \Phi_j \} \right] \\ &= \sum_j \theta_j^* \mathbb{E} \left[\exp \left\{ s \left(\nu(P - P_n) \sum_k \hat{\theta}_k \ell_{e_k}(Z) + \Phi_j + \frac{1}{s} \sum_k \hat{\theta}_k \log(\pi_k) \right) \right\} \right] \\ &= \sum_j \theta_j^* \mathbb{E} \left[\exp \left\{ s \left(\nu(P - P_n) \left(\sum_k \hat{\theta}_k \ell_{e_k}(Z) - \ell_{e_j}(Z) \right) - \mu \sum_k \hat{\theta}_k \|f_k - f_j\|^2 + \frac{1}{s} \sum_k \hat{\theta}_k \log(\pi_k) \right) \right\} \right] \\ &= \sum_j \theta_j^* \mathbb{E} \left[\exp \left\{ s \left(\nu(P - P_n) \sum_k \hat{\theta}_k (\ell_{e_k}(Z) - \ell_{e_j}(Z)) - \mu \sum_k \hat{\theta}_k \|f_k - f_j\|^2 + \frac{1}{s} \sum_k \hat{\theta}_k \log(\pi_k) \right) \right\} \right] \\ &\leq \sum_j \theta_j^* \mathbb{E} \left[\exp \left\{ s \max_{\theta} \left(\nu(P - P_n) \sum_k \theta_k (\ell_{e_k}(Z) - \ell_{e_j}(Z)) - \mu \sum_k \theta_k \|f_k - f_j\|^2 + \frac{1}{s} \sum_k \theta_k \log(\pi_k) \right) \right\} \right] \end{aligned}$$

Since the objective in the exponent is linear in θ it takes its value at one of the extreme points e_1, \dots, e_M :

$$\begin{aligned} I &\leq \sum_j \theta_j^* \mathbb{E} \left[\exp \left\{ s \max_k \left(\nu(P - P_n) (\ell_{e_k}(Z) - \ell_{e_j}(Z)) - \mu \|f_k - f_j\|^2 + \frac{1}{s} \log(\pi_k) \right) \right\} \right] \\ &= \sum_j \theta_j^* \mathbb{E} \left[\max_k \pi_k \exp \{ s (\nu(P - P_n) (\ell_{e_k}(Z) - \ell_{e_j}(Z)) - \mu \|f_k - f_j\|^2) \} \right] \\ &\leq \sum_j \theta_j^* \mathbb{E} \left[\sum_k \pi_k \exp \{ s (\nu(P - P_n) (\ell_{e_k}(Z) - \ell_{e_j}(Z)) - \mu \|f_k - f_j\|^2) \} \right] \\ &\leq \sum_j \theta_j^* \sum_k \pi_k \mathbb{E} \left[\exp \{ s (\nu(P - P_n) (\ell_{e_k}(Z) - \ell_{e_j}(Z)) - \mu \|f_k - f_j\|^2) \} \right] \end{aligned}$$

By Lipschitzness of ℓ , we have that:

$$\|\ell_{e_k} - \ell_{e_j}\|^2 \leq C_b^2 \|f_k - f_j\|^2$$

Thus:

$$\begin{aligned} I &\leq \sum_j \theta_j^* \sum_k \pi_k \mathbb{E} \left[\exp \left\{ s \left(\nu(P - P_n) (\ell_{e_k}(Z) - \ell_{e_j}(Z)) - \frac{\mu}{C_b^2} \|\ell_{e_k} - \ell_{e_j}\|^2 \right) \right\} \right] \\ &= \sum_j \theta_j^* \sum_k \pi_k \mathbb{E} \left[\exp \left\{ s \left(\nu(P - P_n) (\ell_{e_k}(Z) - \ell_{e_j}(Z)) - \frac{\mu}{C_b^2} P(\ell_{e_k}(Z) - \ell_{e_j}(Z))^2 \right) \right\} \right] \\ &= \sum_j \theta_j^* \sum_k \pi_k \mathbb{E} \left[\exp \left\{ s\nu \left((P - P_n) (\ell_{e_k}(Z) - \ell_{e_j}(Z)) - \frac{\mu}{\nu C_b^2} P(\ell_{e_k}(Z) - \ell_{e_j}(Z))^2 \right) \right\} \right] \end{aligned}$$

Invoking Lemma 23 with:

$$\begin{aligned} V_i &= \ell_{e_j}(Z_i) - \ell_{e_k}(Z_i) \\ c_0 &= \frac{\mu}{\nu C_b^2} \\ \lambda &= \frac{s\nu}{n} \\ c &= 2bC_b \end{aligned}$$

Hence we need:

$$\lambda := \frac{s\nu}{n} \leq \frac{2c_0}{1 + 2c_0c} =: \frac{2\mu}{\nu C_b^2 (1 + 2\frac{\mu}{\nu C_b^2} 2bC_b)} = \frac{2\mu}{(\nu C_b^2 + 4\mu bC_b)} \implies s \leq \frac{2\mu n}{\nu C_b (\nu C_b + 4\mu b)}$$

□

E.2 Putting it all together

With the choice of β and μ as described in the theorem statement, we can invoke the above Lemmas. Letting $C_b(\hat{g})$ the Lipschitz constant associated with an estimate \hat{g} , we have that w.p. $1 - \delta$:

$$R(\hat{\theta}) \leq \min_j \left(R(e_j) + \frac{\beta}{n} \log(1/\pi_j) \right) + O \left(\frac{\log(1/\delta) \max\{\frac{1}{\mu} C_b(\hat{g})^2, b C_b(\hat{g})\}}{n} \right) + \frac{2}{\mu} \text{ERROR}(\hat{g})^2$$

E.3 Basic concentration proofs

Lemma 21. For any random variable Z , if $\mathbb{E}[\exp\{Z\}] \leq 1$, then $\Pr(Z \geq \log(1/\delta)) \leq \delta$.

Proof. Suppose that one shows that:

$$\mathbb{E}[\exp\{Z\}] \leq 1$$

Then by Markov's inequality:

$$\Pr(Z \geq x) = \Pr(\exp(Z) \geq \exp(x)) \leq \frac{\mathbb{E}[\exp(Z)]}{\exp(x)} \leq \frac{1}{\exp(x)} = \exp(-x)$$

Thus:

$$\Pr(Z \geq \log(1/\delta)) \leq \delta$$

□

Lemma 22. *Suppose $\epsilon_1, \dots, \epsilon_n$ are independent Rademacher random variables drawn equiprobably in $\{-1, 1\}$ and that a_1, \dots, a_n are real numbers. Then:*

$$\mathbb{E}_{\epsilon_{1:n}} \left[\exp \left\{ \frac{1}{n} \sum_i \epsilon_i a_i \right\} \right] \leq \exp \left\{ \frac{1}{2n^2} \sum_i a_i^2 \right\}$$

Proof. First note that for by the derivation in Example 2.2 of https://www.stat.berkeley.edu/~mjwain/stat210b/Chap2_TailBounds_Jan22_2015.pdf, we have that:

$$\mathbb{E}[\exp(\epsilon_i a_i/n)] \leq \exp(a_i^2/2n^2)$$

Invoking now independence of the ϵ_i we have:

$$\begin{aligned} \mathbb{E}_{\epsilon_{1:n}} \left[\exp \left\{ \frac{1}{n} \sum_i \epsilon_i a_i \right\} \right] &= \mathbb{E}_{\epsilon_{1:n}} \left[\prod_{i=1}^n \exp \{ \epsilon_i a_i/n \} \right] \\ &= \prod_{i=1}^n \mathbb{E}_{\epsilon_{1:n}} [\exp \{ \epsilon_i a_i/n \}] && \text{(independence)} \\ &\leq \prod_{i=1}^n \exp \left(\frac{a_i^2}{2n^2} \right) = \exp \left\{ \frac{1}{2n^2} \sum_i a_i^2 \right\} \end{aligned}$$

□

The following proposition can also be proven as an application of Bernstein's inequality.

Lemma 23 (Proposition 1, Lecué and Rigollet (2014)). *Let V_1, \dots, V_n be i.i.d. random variables, such that $|V_i| \leq c$ a.s.. Let $c_0 > 0$, then for any $0 < \lambda < \frac{2c_0}{1+2c_0c}$:*

$$\mathbb{E} \left[\exp \left\{ n\lambda \left(\frac{1}{n} \sum_i (V_i - \mathbb{E}[V_i] - c_0 \mathbb{E}[V_i^2]) \right) \right\} \right] \leq 1$$

E.4 Auxiliary Lemmas

Lemma 24 (Symmetrization Inequality (Lecué and Rigollet, 2014)). *let \mathcal{F} be a function class, $A_f, f \in \mathcal{F}$ be a given function on \mathcal{F} and Φ be a convex non-decreasing function then*

$$\mathbb{E}\Phi \left(\sup_{f \in \mathcal{F}} \left[\mathbb{E}[f(Z)] - \frac{1}{n} \sum_{i=1}^n f(Z_i) - A_f \right] \right) \leq \mathbb{E}\Phi \left(2 \sup_{f \in \mathcal{F}} \left[\frac{1}{n} \sum_{i=1}^n \epsilon_i f(Z_i) - A_f \right] \right)$$

Lemma 25 (Contraction Inequality, Theorem 4.12 in (Ledoux and Talagrand, 1991)). *Let $F : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be convex and increasing. Let further $\varphi_i : \mathbb{R} \rightarrow \mathbb{R}, i \leq N$, be such that $|\varphi_i(s) - \varphi_i(t)| \leq |s - t|$ and $\varphi_i(0) = 0$. Then, for any bounded subset T in \mathbb{R}^N*

$$\mathbb{E}F \left(\frac{1}{2} \sup_T \sum_{i=1}^N \epsilon_i \varphi_i(t_i) \right) \leq \mathbb{E}F \left(\sup_T \sum_{i=1}^N \epsilon_i t_i \right)$$

F PROOF OF COROLLARIES FOR SQUARE LOSSES

Proof of Corollary 4. Note that if we assume that $\hat{Y}(\hat{g})$ is uniformly bounded for any $\hat{g} \in G$ with $d(\hat{g}, g_0) \leq \epsilon$ for some constant ϵ and since for n larger than the designated constant, we have that w.p. $1 - \delta$, $\hat{Y}(\hat{g})$ is absolutely and uniformly bounded by U . Moreover, $f_j(X)$ is absolutely and uniformly bounded by U . Note that by an exact first order Taylor expansion:

$$\begin{aligned} |\ell(Z; f(X), g) - \ell(Z; f_*(X), g)| &= \left| \int_0^1 2(\hat{Y}(g) - (1 - \tau)f(X) - \tau f_*(X))(f(X) - f_*(X)) d\tau \right| \\ &\leq 2 \left(\sup_{f \in F} |f(X)| + |\hat{Y}(g)| \right) |f(X) - f_*(X)| \\ &\leq 2U |f(X) - f_*(X)| \end{aligned}$$

Thus the loss function is $C_b = 2U$ lipschitz and absolutely bounded by $4U^2$. Moreover, we have that the loss function $\ell(Z; f(X), g_0)$ is 2-strongly convex in f in expectation. Thus, we can take $\sigma = 2$ in Theorem 3. Thus we can apply the main theorem with $\nu = 1/2$, $\mu = 1/14$ and $\beta = \max\{112U^2, 56U^3\}$.

Finally, we analyze the quantity $\mathcal{E}(\hat{g})$. We first note that for any θ, g , we have:

$$\begin{aligned}\ell_{\theta,g}(Z) - \ell_{\theta^*,g}(Z) &= (\hat{Y}(g) - f_\theta(X))^2 - (\hat{Y}(g) - f_{\theta^*}(X))^2 \\ &= f_\theta(X)^2 - f_{\theta^*}(X)^2 + 2\hat{Y}(g)(f_{\theta^*}(X) - f_\theta(X))\end{aligned}$$

Moreover, for any $\theta \in \Theta$:

$$\begin{aligned}\sum_j \theta_j \ell_{e_j,g}(Z) &= \sum_j \theta_j (\hat{Y}(g) - f_j(X))^2 \\ &= \hat{Y}(g)^2 \sum_j \theta_j + \sum_j \theta_j f_j(X)^2 - 2 \sum_j \theta_j \hat{Y}(g) f_j(X) \\ &= \hat{Y}(g)^2 + \sum_j \theta_j f_j(X)^2 - 2 \sum_j \theta_j \hat{Y}(g) f_j(X) \quad (\sum_j \theta_j = 1) \\ &= \hat{Y}(g)^2 + \sum_j \theta_j f_j(X)^2 - 2\hat{Y}(g) f_\theta \quad (f_\theta = \sum_j \theta_j f_j)\end{aligned}$$

Thus we have that:

$$\sum_j \theta_j \ell_{e_j,g}(Z) - \sum_j \theta_j^* \ell_{e_j,g}(Z) = \sum_j (\theta_j - \theta_j^*) f_j(X)^2 + 2\hat{Y}(g)(f_{\theta^*}(X) - f_\theta(X))$$

Then by the definition of $\tilde{\ell}_{\theta,g}$, we have:

$$\ell_{\theta,g}(Z) - \ell_{\theta^*,g}(Z) = (1 - \nu)(f_\theta(X)^2 - f_{\theta^*}(X)^2) + \nu \sum_j (\theta_j - \theta_j^*) f_j(X)^2 + 2\hat{Y}(g)(f_{\theta^*}(X) - f_\theta(X))$$

Only the last part depends on the parameter g . Thus we get:

$$\mathcal{E}(\hat{g}) = 2\mathbb{E}[(\hat{Y}(g_0) - \hat{Y}(\hat{g}))(f_{\theta^*}(X) - f_{\hat{\theta}}(X))] = 2\mathbb{E}\left[\mathbb{E}\left[\hat{Y}(g_0) - \hat{Y}(\hat{g}) \mid X\right] (f_{\theta^*}(X) - f_{\hat{\theta}}(X))\right]$$

Applying a Cauchy-Schwarz inequality, we have:

$$\mathcal{E}(\hat{g}) = 2\sqrt{\mathbb{E}\left[\mathbb{E}\left[\hat{Y}(g_0) - \hat{Y}(\hat{g}) \mid X\right]^2\right]} \|f_{\theta^*} - f_{\hat{\theta}}\|_{L^2}$$

and we can take:

$$\text{ERROR}(\hat{g}) = 2\sqrt{\mathbb{E}\left[\mathbb{E}\left[\hat{Y}(g_0) - \hat{Y}(\hat{g}) \mid X\right]^2\right]}$$

in Theorem 3. Finally, note that since $\mathbb{E}[\hat{Y}(g_0) \mid X] = \tau_0(X)$, the excess risk evaluated at the true parameters satisfies that:

$$\begin{aligned}R(\theta, g_0) - R(\theta', g_0) &= \mathbb{E}\left[f_\theta(X)^2 - f_{\theta'}(X)^2 + 2\hat{Y}(g_0)(f_\theta(X) - f_{\theta'}(X))\right] \\ &= \mathbb{E}\left[f_\theta(X)^2 - f_{\theta'}(X)^2 + 2\tau_0(X)(f_\theta(X) - f_{\theta'}(X))\right] \\ &= \mathbb{E}\left[(f_\theta(X) - \tau_0(X))^2 - (f_{\theta'}(X) - \tau_0(X))^2\right] \\ &= \|f_\theta - \tau_0\|^2 - \|f_{\theta'} - \tau_0\|^2\end{aligned}$$

Thus we have that:

$$R(\hat{\theta}, g_0) - \min_j R(e_j, g_0) = \|f_{\hat{\theta}} - \tau_0\|^2 - \min_j \|f_j - \tau_0\|^2$$

□

Proof of Corollary 5. The proof is identical to Corollary 4, but we simply replace the application of the Cauchy-Schwarz inequality with the inequality:

$$\begin{aligned} 2\mathbb{E} \left[\mathbb{E} \left[\hat{Y}(g_0) - \hat{Y}(\hat{g}) \mid X \right] (f_{\theta^*}(X) - f_{\hat{\theta}}(X)) \right] &\leq 2\mathbb{E} \left[\left| \mathbb{E} \left[\hat{Y}(g_0) - \hat{Y}(\hat{g}) \mid X \right] \right| \right] \|f_{\theta^*} - f_{\hat{\theta}}\|_{L^\infty} \\ &\leq 2C \mathbb{E} \left[\left| \mathbb{E} \left[\hat{Y}(g_0) - \hat{Y}(\hat{g}) \mid X \right] \right| \right] \|f_{\theta^*} - f_{\hat{\theta}}\|_{L^2}^\gamma \end{aligned}$$

□

F.1 Relaxed Nuisance Guarantees Under Eigenvalue Condition

Lemma 26. *Let F be the random column vector defined as $F := (f_1(X), \dots, f_M(X))$. Suppose that $\mathbb{E}[FF'] \succeq cI$ and $\|F\|_{L^\infty} \leq U$. Then we can take $\gamma = 1$ and $C = 2U\sqrt{\frac{M}{c}}$ in Corollary 5.*

Here, the number of models M and not the logarithm of them appears. However, assuming that the nuisance error term decays faster than $n^{-1/2}$, this linear in M term will only appear in the lower-order terms.

Proof. Let F be the random column vector defined as $F := (f_1(X), \dots, f_M(X))$ and assume that $\|F\|_{L^\infty} \leq U$. Noting that:

$$f_{\theta^*}(X) - f_{\hat{\theta}}(X) = F'(\theta^* - \hat{\theta}).$$

we have:

$$\begin{aligned} \|f_{\theta^*}(X) - f_{\hat{\theta}}(X)\|_{L^\infty} &\leq \|F\|_{L^\infty} \|\theta^* - \hat{\theta}\|_1 \\ &\leq U \|\theta^* - \hat{\theta}\|_1 \leq U\sqrt{M} \|\theta^* - \hat{\theta}\|_2 \end{aligned}$$

Then we if we further assume that $\mathbb{E}[FF'] \succeq cI$, i.e. the CATE models all contain independent signal components and are not colinear. Then we can write:

$$\begin{aligned} \|f_{\theta^*}(X) - f_{\hat{\theta}}(X)\|_{L^\infty} &\leq U\sqrt{M} \sqrt{\|\theta^* - \hat{\theta}\|_2^2} \\ &\leq U\sqrt{M} \sqrt{\frac{1}{c}(\theta^* - \hat{\theta})\mathbb{E}[FF'](\theta^* - \hat{\theta})} \\ &= U\sqrt{M} \sqrt{\frac{1}{c}\mathbb{E}[(F'(\theta^* - \hat{\theta}))^2]} \\ &= U\sqrt{M} \frac{1}{\sqrt{c}} \|f_{\theta^*} - f_{\hat{\theta}}\|_{L^2} \end{aligned}$$

Thus under these further assumptions, we can set:

$$C = 2U\sqrt{\frac{M}{c}}$$

in Corollary 5. □

G PROOF OF MAIN THEOREM WITHOUT SAMPLE SPLITTING

We will solely change the statement and proof of Lemma 18 and subsequently Lemma 19 and Lemma 20 will be invoked at g_0 and not \hat{g} . The rest of the Lemmas do not depend on whether \hat{g} was trained on the same samples or not as the ones that were used to construct $\hat{\theta}$.

Lemma 27. Define the centered and offset empirical processes:

$$\begin{aligned}
 Z_{n,1} &= (1 - \nu)(P - P_n)(\ell_{\hat{\theta},g_0}(Z) - \ell_{\theta^*,g_0}(Z)) - \mu \sum_j \hat{\theta}_j \|f_j - f_{\theta^*}\|^2 - \frac{1}{s}K(\hat{\theta}) \\
 Z_{n,2} &= \nu(P - P_n) \sum_j (\hat{\theta}_j - \theta_j^*) \ell_{e_j,g_0}(Z) - \mu \hat{\theta}^\top H \theta^* - \frac{1}{s}K(\hat{\theta}) & H_{ij} &= \|f_i - f_j\|^2 \\
 Z_{n,3} &= \max_j \left[(P - P_n)(f_j(X) - f_{\theta^*}(X))^2 - \mu \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \log(\pi_j) \right] \\
 Z_{n,4} &= (P - P_n)(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2
 \end{aligned}$$

Assume that:

$$\mathcal{E}(\hat{g}) := P(\tilde{\ell}_{\hat{\theta},g_0}(Z) - \tilde{\ell}_{\theta^*,g_0}(Z)) - P(\tilde{\ell}_{\hat{\theta},\hat{g}}(Z) - \tilde{\ell}_{\theta^*,\hat{g}}(Z)) \leq \text{ERROR}(\hat{g}) \|f_\theta - f_{\theta^*}\|$$

and $\ell(f, g_0)$ is 1-strongly convex in expectation with respect to f since we consider square losses. For any $\mu \leq \min\{\nu, 1 - \nu\} \frac{1}{14}$:

$$R(\hat{\theta}) \leq \min_j R(e_j) + 2(Z_{n,1} + Z_{n,2} + Z_{n,3} + Z_{n,4}) + \frac{1}{\mu} \text{ERROR}(\hat{g})^2$$

Proof.

$$\begin{aligned}
 \tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0) &= P(\tilde{\ell}_{\hat{\theta},g_0}(Z) - \tilde{\ell}_{\theta^*,g_0}(Z)) \\
 &= P(\tilde{\ell}_{\hat{\theta},\hat{g}}(Z) - \tilde{\ell}_{\theta^*,\hat{g}}(Z)) + \mathcal{E}(\hat{g}) \\
 &\leq (P - P_n)(\tilde{\ell}_{\hat{\theta},\hat{g}}(Z) - \tilde{\ell}_{\theta^*,\hat{g}}(Z)) - \frac{\beta}{n}(K(\hat{\theta}) - K(\theta^*)) + \mathcal{E}(\hat{g}) \\
 &\quad \text{(By optimality of } \hat{\theta}, P_n(\tilde{\ell}_{\hat{\theta},\hat{g}}(Z) - \tilde{\ell}_{\theta^*,\hat{g}}(Z)) + \frac{\beta}{n}(K(\hat{\theta}) - K(\theta^*)) \leq 0) \\
 &= (P - P_n)(\tilde{\ell}_{\hat{\theta},g_0}(Z) - \tilde{\ell}_{\theta^*,g_0}(Z)) + \Delta(\hat{g}, \hat{\theta}) - \frac{\beta}{n}(K(\hat{\theta}) - K(\theta^*)) + \mathcal{E}(\hat{g}) \\
 &= \mu \sum_j \hat{\theta}_j \|f_j - f_{\theta^*}\|^2 + \mu \hat{\theta}^\top H \theta^* + Z_{n,1} + Z_{n,2} \\
 &\quad + \frac{2}{s}K(\hat{\theta}) - \frac{\beta}{n}(K(\hat{\theta}) - K(\theta^*)) + \Delta(\hat{g}, \hat{\theta}) + \mathcal{E}(\hat{g}) \quad \text{(definition of } Z_{n,1}, Z_{n,2})
 \end{aligned}$$

where:

$$\Delta(\hat{g}, \hat{\theta}) = (P - P_n) \left\{ \tilde{\ell}_{\hat{\theta},\hat{g}}(Z) - \tilde{\ell}_{\theta^*,\hat{g}}(Z) - \left(\tilde{\ell}_{\hat{\theta},g_0}(Z) - \tilde{\ell}_{\theta^*,g_0}(Z) \right) \right\}$$

For the case of square losses, the above difference in differences simplifies to:

$$\begin{aligned}
 \Delta(\hat{g}, \hat{\theta}) &= (P - P_n) \left\{ 2(\hat{Y}(\hat{g}) - \hat{Y}(g_0))(f_{\hat{\theta}}(X) - f_{\theta^*}(X)) \right\} \\
 &= \sum_j \hat{\theta}_j \left[(P - P_n) \left\{ 2(\hat{Y}(\hat{g}) - \hat{Y}(g_0))(f_j(X) - f_{\theta^*}(X)) \right\} - \frac{1}{s} \log\left(\frac{1}{\pi_j}\right) \right] + \frac{1}{s}K(\hat{\theta}) \\
 &= \mu \sum_j \hat{\theta}_j \|f_j - f_{\hat{\theta}^*}\|^2 + \frac{1}{s}K(\hat{\theta}) \\
 &\quad + \sum_j \hat{\theta}_j \left[(P - P_n) \left\{ 2(\hat{Y}(\hat{g}) - \hat{Y}(g_0))(f_j(X) - f_{\theta^*}(X)) \right\} - \mu \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \log(\pi_j) \right] \\
 &= \mu \sum_j \hat{\theta}_j \|f_j - f_{\hat{\theta}^*}\|^2 + \frac{1}{s}K(\hat{\theta}) \\
 &\quad + \max_{\theta \in \Theta} \sum_j \theta_j \left[(P - P_n) \left\{ 2(\hat{Y}(\hat{g}) - \hat{Y}(g_0))(f_j(X) - f_{\theta^*}(X)) \right\} - \mu \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \log(\pi_j) \right]
 \end{aligned}$$

Since the last term is linear in $\hat{\theta}$ and Θ is the simplex, it is maximized at a corner, which yields:

$$\begin{aligned}
 \Delta(\hat{g}, \hat{\theta}) &\leq \mu \sum_j \hat{\theta}_j \|f_j - f_{\hat{\theta}^*}\|^2 + \frac{1}{s} K(\hat{\theta}) \\
 &\quad + \max_j \left[(P - P_n) \left\{ 2(\hat{Y}(\hat{g}) - \hat{Y}(g_0)) (f_j(X) - f_{\theta^*}(X)) \right\} - \mu \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \log(\pi_j) \right] \\
 &\leq \mu \sum_j \hat{\theta}_j \|f_j - f_{\hat{\theta}^*}\|^2 + \frac{1}{s} K(\hat{\theta}) + (P - P_n) (\hat{Y}(\hat{g}) - \hat{Y}(g_0))^2 \\
 &\quad + \max_j \left[(P - P_n) (f_j(X) - f_{\theta^*}(X))^2 - \mu \|f_j - f_{\theta^*}\|^2 + \frac{1}{s} \log(\pi_j) \right] \\
 &\leq \mu \sum_j \hat{\theta}_j \|f_j - f_{\hat{\theta}^*}\|^2 + Z_{n,3} + Z_{n,4} + \frac{1}{s} K(\hat{\theta})
 \end{aligned}$$

By simple algebraic manipulations note that for any $\theta \in \Theta$:

$$\begin{aligned}
 \sum_j \hat{\theta}_j \|f_j - f_{\theta^*}\|^2 &= V(\hat{\theta}) + \|f_{\hat{\theta}} - f_{\theta^*}\|^2 \\
 \hat{\theta}^\top H \theta^* &= V(\hat{\theta}) + V(\theta^*) + \|f_{\hat{\theta}} - f_{\theta^*}\|^2
 \end{aligned}$$

Letting $Z_n = Z_{n,1} + Z_{n,2} + Z_{n,3} + Z_{n,4}$, we have:

$$\begin{aligned}
 \tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0) &\leq 3\mu V(\hat{\theta}) + \mu V(\theta^*) + 3\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 + Z_n + \frac{3}{s} K(\hat{\theta}) - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + \mathcal{E}(\hat{g}) \\
 &\leq 3\mu V(\hat{\theta}) + \mu V(\theta^*) + 3\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 + Z_n \\
 &\quad + \frac{3}{s} K(\hat{\theta}) - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + \text{ERROR}(\hat{g}) \|f_{\hat{\theta}} - f_{\theta^*}\| \\
 &\leq 3\mu V(\hat{\theta}) + \mu V(\theta^*) + 4\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 + Z_n \\
 &\quad + \frac{3}{s} K(\hat{\theta}) - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + \frac{1}{\mu} \text{ERROR}(\hat{g})^2 \tag{6}
 \end{aligned}$$

where the last inequality follows by an application of the AM-GM inequality. By Lemma 17, we have:

$$\begin{aligned}
 \tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0) &\geq \mu(V(\theta^*) - V(\hat{\theta})) - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + \left(\frac{(1-\nu)\sigma}{2} - \mu \right) \|f_{\hat{\theta}} - f_{\theta^*}\|^2 \\
 &\geq \mu(V(\theta^*) - V(\hat{\theta})) - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + 8\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 \\
 &\quad \text{(Since: } \sigma = 2 \text{ and } \mu \leq \frac{(1-\nu)}{18} \implies \frac{(1-\nu)\sigma}{2} \geq 9\mu)
 \end{aligned}$$

Combining the latter upper and lower bound on $\tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0)$, we have:

$$4\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2 \leq 4\mu V(\hat{\theta}) + Z_n + \frac{3}{s} K(\hat{\theta}) + \frac{1}{\mu} \text{ERROR}(\hat{g})^2$$

Plugging the bound on $4\mu \|f_{\hat{\theta}} - f_{\theta^*}\|^2$ back into Equation (6) we have:

$$\tilde{R}(\hat{\theta}, g_0) - \tilde{R}(\theta^*, g_0) \leq 7\mu V(\hat{\theta}) + \mu V(\theta^*) + 2Z_n - \frac{\beta}{n} (K(\hat{\theta}) - K(\theta^*)) + \frac{6}{s} K(\hat{\theta}) + \frac{2}{\mu} \text{ERROR}(\hat{g})^2$$

By Lemma 16, we have:

$$\begin{aligned}
 R(\hat{\theta}, g_0) &\leq \tilde{R}(\theta^*, g_0) + \mu V(\theta^*) + \left(7\mu - \frac{\nu\sigma}{2}\right) V(\hat{\theta}) + 2Z_n + \frac{2}{\mu} \text{ERROR}(\hat{g})^2 \\
 &\quad + \left(\frac{6}{s} - \frac{\beta}{n}\right) K(\hat{\theta}) + \frac{\beta}{n} K(\theta^*) \\
 &\leq \tilde{R}(\theta^*, g_0) + \mu V(\theta^*) + \frac{\beta}{n} K(\theta^*) + 2Z_n + \frac{2}{\mu} \text{ERROR}(\hat{g})^2 \\
 &\quad \text{(since, } \sigma = 2, \mu \leq \frac{\nu}{14} \text{ and } V(\hat{\theta}) > 0 \text{ and } \beta \geq \frac{6n}{s} \text{ and } K(\hat{\theta}) \geq 0) \\
 &= \min_{\theta \in \Theta} \left[\tilde{R}(\theta, g_0) + \mu V(\theta) + \frac{\beta}{n} K(\theta^*) \right] + 2Z_n + \frac{2}{\mu} \text{ERROR}(\hat{g})^2 \quad \text{(Since } \theta_* \text{ minimizes } H(\theta, g_0)) \\
 &\leq \min_j \left[\tilde{R}(e_j, g_0) + \mu V(e_j) + \frac{\beta}{n} \log\left(\frac{1}{\pi_j}\right) \right] + 2Z_n + \frac{2}{\mu} \text{ERROR}(\hat{g})^2 \\
 &= \min_j \left[R(e_j, g_0) + \frac{\beta}{n} \log\left(\frac{1}{\pi_j}\right) \right] + 2Z_n + \frac{2}{\mu} \text{ERROR}(\hat{g})^2 \\
 &\quad \text{(Since } \tilde{R}(e_j, g_0) = R(e_j, g_0) \text{ and } V(e_j) = 0 \text{ for all } e_j)
 \end{aligned}$$

□

Lemma 28. Assuming that $f \in F$ are uniformly and absolutely and bounded in $[-b, b]$, for any $s \leq \frac{n\mu}{2b^2(1+\mu)}$:

$$\Pr\left(Z_{n,3} \geq \frac{\log(1/\delta)}{s}\right) \leq \delta$$

Proof. It suffices to show that:

$$I := \mathbb{E}[\exp sZ_{n,3}] \leq 1$$

Note first that:

$$\begin{aligned}
 \mathbb{E}[\exp sZ_{n,3}] &= \mathbb{E}\left[\exp\left\{s \max_j \left[(P - P_n)(f_j(Z) - f_{\theta^*}(Z))^2 - \mu P(f_j(Z) - f_{\theta^*}(Z))^2 + \frac{1}{s} \log(\pi_j) \right] \right\}\right] \\
 &\leq \mathbb{E}\left[\max_j \exp\left\{s \left((P - P_n)(f_j(Z) - f_{\theta^*}(Z))^2 - \mu P(f_j(Z) - f_{\theta^*}(Z))^2 + \frac{1}{s} \log(\pi_j) \right) \right\}\right] \\
 &\leq \sum_j \pi_j \mathbb{E}\left[\exp\left\{s \left((P - P_n)(f_j(Z) - f_{\theta^*}(Z))^2 - \mu P(f_j(Z) - f_{\theta^*}(Z))^2 \right) \right\}\right] \\
 &\leq \max_j \mathbb{E}\left[\exp\left\{s \left((P - P_n)(f_j(Z) - f_{\theta^*}(Z))^2 - \mu P(f_j(Z) - f_{\theta^*}(Z))^2 \right) \right\}\right]
 \end{aligned}$$

Since $f_j(Z) \in [-b, b]$, we have that $f_j(Z) - f_{\theta^*}(Z) \in [-2b, 2b]$. Thus:

$$(f_j(Z) - f_{\theta^*}(Z))^2 \geq \frac{(f_j(Z) - f_{\theta^*}(Z))^2}{4b^2} (f_j(Z) - f_{\theta^*}(Z))^2 = \frac{1}{4b^2} (f_j(Z) - f_{\theta^*}(Z))^4.$$

Thus:

$$I \leq \max_j \mathbb{E}_{Z_{1:n}} \left[\exp\left\{s \left((P_n - P)(f_j(Z) - f_{\theta^*}(Z))^2 - \frac{\mu}{4b^2} P(f_j(X) - f_{\theta^*}(X))^4 \right) \right\}\right]$$

It suffices to show that each term j is upper bounded by 1.

Applying Lemma 23 with:

$$\begin{aligned}
 V_i &= (f_j(Z_i) - f_{\theta^*}(Z_i))^2 \\
 c &= 4b^2 \\
 \lambda &= s \\
 c_0 &= \frac{\mu}{4b^2}
 \end{aligned}$$

we have that as long as:

$$\lambda := \frac{s}{n} \leq \frac{2c_0}{1+2c_0c} =: \frac{2 \cdot \mu/4b^2}{1+\mu} = \frac{\mu}{2b^2(1+\mu)}$$

Then:

$$\mathbb{E}_{Z_{1:n}} \left[\exp \left\{ s \left((P_n - P)(f_j(Z) - f_{\theta^*}(Z))^2 - \frac{\mu}{4b^2} P(f_j(X) - f_{\theta^*}(X))^4 \right) \right\} \right] \leq 1$$

Thus we need that:

$$s \leq \frac{n\mu}{2b^2(1+\mu)}$$

which is satisfied by our assumptions on s . \square

To bound $Z_{n,4}$ we will use the following concentration inequality of based on critical radius (e.g. Wainwright (2019)). See (Foster and Syrgkanis, 2023) for a proof of this theorem:

Lemma 29 (Localized Concentration, (Foster and Syrgkanis, 2023)). *For any $h \in \mathcal{H} := \prod_{i=1}^d \mathcal{H}_i$ be a multi-valued outcome function, that is almost surely absolutely bounded by a constant. Let $\ell(Z; h(W)) \in \mathbb{R}$ be a loss function that is L -Lipschitz in $h(W)$, with respect to the ℓ_2 norm. Fix any $h_* \in \mathcal{H}$ and let $\delta_n = \Omega \left(\sqrt{\frac{d \log \log(n) + \log(1/\zeta)}{n}} \right)$ be an upper bound on the critical radius of $\text{star}(\mathcal{H}_i - h_{i,*})$ for $i \in [d]$. Then w.p. $1 - \zeta$:*

$$\forall h \in \mathcal{H} : |(P_n - P)[\ell(Z; h(W)) - \ell(Z; h_*(W))]| = 18L \left(d\delta_n \sum_{i=1}^d \|h_i - h_{i,*}\|_{L_2} + d\delta_n^2 \right)$$

If the loss is linear in $h(W)$, i.e. $\ell(Z; h(W) + h'(W)) = \ell(Z; h(W)) + \ell(Z; h'(W))$ and $\ell(Z; \alpha h(W)) = \alpha \ell(Z; h(W))$ for any scalar α , then it suffices that we take $\delta_n = \Omega \left(\sqrt{\frac{\log(1/\zeta)}{n}} \right)$ that upper bounds the critical radius of $\text{star}(\mathcal{H}_i - h_{i,*})$ for $i \in [d]$.

Lemma 30. *Let r_n denote the critical radius of the function space $\text{star}(Y_G)$ where $Y_G := \{(\hat{Y}(g) - \hat{Y}(g_0))^2 | g \in G\}$ and assume that $\hat{Y}(g)$ is uniformly bounded by U for all $g \in G$. Then for $\delta_n \geq \max \left\{ r_n, \Omega \left(\sqrt{\frac{\log(1/\delta)}{n}} \right) \right\}$, w.p. $1 - \delta$*

$$Z_{n,4} \leq 18 \left(\delta_n \sqrt{\mathbb{E} \left[(\hat{Y}(g) - \hat{Y}(g_0))^4 \right]} + \delta_n^2 \right) \quad (7)$$

Proof. We can apply Lemma 29 with:

$$h_g(W) := (\hat{Y}(g) - \hat{Y}(g_0))^2, g \in G \quad h_*(W) := 0 \quad \ell(Z; h(W)) := h(W)$$

where W is assumed to contain all the random variables that go into the calculation of the proxy labels $\hat{Y}(g)$. Note that the loss function is 1-Lipschitz in h and we have a scalar-valued function h , i.e. $d = 1$ and the loss function is linear in h . Thus we get that if $\delta_n = \Omega \left(\sqrt{\frac{\log(1/\zeta)}{n}} \right)$ upper bounds the critical radius r_n of $\text{star}(Y_G)$, then w.p. $1 - \zeta$:

$$|(P_n - P)h_{\hat{g}}(W)| = 18 (\delta_n \|h_{\hat{g}}\|_{L_2} + \delta_n^2) \leq 18 \left(\delta_n \sqrt{\mathbb{E} \left[(\hat{Y}(g) - \hat{Y}(g_0))^4 \right]} + \delta_n^2 \right)$$

\square

G.1 Putting it all together

Applying the AM-GM inequality to Equation (7), we get:

$$Z_{n,4} \leq \mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^4] + 23\delta_n^2 = \mathbb{E}[(\hat{Y}(\hat{g}) - \hat{Y}(g_0))^4] + O\left(r_n^2 + \frac{\log(1/\delta)}{n}\right)$$

Assuming $\hat{Y}(g)$ be uniformly bounded by U and $f_j(X)$ is absolutely and uniformly bounded by U , the square loss is $4U$ -Lipschitz. Choosing $\nu = 1/2$ and combining all the above Lemmas, we have that w.p. $1 - \delta$:

$$R(\hat{\theta}) \leq \min_j \left(R(e_j) + \frac{\beta}{n} \log(1/\pi_j) \right) + O\left(r_n^2 + \frac{\log(\log(n)) + \log(1/\delta) \max\{\frac{1}{\mu}U^2, bU\}}{n}\right) + \frac{2}{\mu} \text{ERROR}(\hat{g})^2$$

H ADDITIONAL REMARKS FOR CATE WITH INSTRUMENTS

Remark 3 (Beyond Stratified Trials with Non-Compliance). *When the instrument policy $\pi_0(X)$ is un-known, i.e. when we have a conditional instrument that corresponds to some natural experiment, then the quantity $\hat{Y}(\hat{g})$ is sensitive to errors in the estimate $\hat{\pi}$ of π_0 . To reduce this sensitivity, we can incorporate residualization of Y and D (as also proposed in Syrgkanis et al. (2019)):*

$$\hat{Y}(g) = \hat{\tau}(X) + \frac{(\tilde{Y} - \hat{\tau}(X)\tilde{D})\tilde{Z}}{\hat{\beta}(X)}$$

where $\tilde{Y} = Y - \hat{h}(X)$, $\tilde{D} = D - \hat{r}(X)$ and $\tilde{Z} = Z - \hat{\pi}(X)$ and \hat{h} is an estimate of $h_0(X) = \mathbb{E}[Y | X]$, \hat{r} is an estimate of $r_0(X) = \mathbb{E}[D | X]$ and $\hat{\pi}$ is an estimate of $\pi_0(X) = \mathbb{E}[Z | X]$. With such a proxy label the error term in the Corollaries presented in this section, will contain extra terms that can be upper bounded by terms of the order $\|\hat{\pi} - \pi_0\|_{L^4}^2 \|\hat{r} - r_0\|_{L^4}^2$, $\|\hat{\pi} - \pi_0\|_{L^4}^2 \|\hat{h} - h_0\|_{L^4}^2$, $\|\hat{\pi} - \pi_0\|_{L^4}^2 \|\hat{\beta} - \beta_0\|_{L^4}^2$, $\|\hat{r} - r_0\|_{L^4}^2 \|\hat{\beta} - \beta_0\|_{L^4}^2$ and $\|\hat{h} - h_0\|_{L^4}^2 \|\hat{\beta} - \beta_0\|_{L^4}^2$.

I GREEDY Q-AGGREGATION FOR SQUARE LOSSES

In this section, we show that we can get comparable theoretical guarantees when using a k-step greedy approximation of the Q-aggregation algorithm.

I.1 Preliminaries

Consider the following k-step greedy algorithm to find an approximate minimizer of a functional $Q(\theta)$, proposed in Dai et al. (2012):

Algorithm 1 Greedy model averaging GMA-0₊

Let $\theta^{(0)} = 0$, $f_{\theta^{(0)}} = 0$
for $k = 1, 2, \dots$ **do**
 Set $a_k = \frac{2}{k+1}$
 $J^{(k)} = \arg \min_j Q((1 - a_k)\theta^{(k-1)} + a_k e_j)$
 $\theta^{(k)} = \arg \min_{\theta \in \Theta} Q(\theta)$ s.t. $\theta_j = 0$ for $j \notin \{J^{(1)}, \dots, J^{(k)}\}$
end for

Theorem 31 (Dai et al. (2012)). *Suppose that $\ell_{\theta, \hat{g}}(Z)$ is of the form $(Y - f_{\theta}(X))^2$ for some label $Y \equiv \hat{Y}(\hat{g})$. Denote with:*

$$Q(\theta) := P_n \tilde{\ell}_{\theta, \hat{g}}(Z) + \frac{\beta}{n} K(\theta) \qquad V_n(\theta) := \sum_{j=1}^M \theta_j P_n (f_j(X) - f_{\theta}(X))^2$$

Then the estimator $f_{\theta^{(k)}}$, where $\theta^{(k)}$ is the output of GMA-0₊ after k steps satisfies:

$$Q(\theta^{(k)}) \leq \min_{\theta \in \Theta} \{Q(\theta) + C_Q V_n(\theta)\}$$

where $C_Q = \frac{4(1-\nu)}{k+3}$.

Proof. We first introduce some short-hand notation that will be convenient for the proof. Let $x_\theta = (f_\theta(X_1), \dots, f_\theta(X_n)) \in \mathbb{R}^n$, denote the vector of values of the function f_θ on the n samples and $y = (Y_1, \dots, Y_n) \in \mathbb{R}^n$ denote the vector of samples. For any vectors $x_\theta, x_{\bar{\theta}} \in \mathbb{R}^n$ define the empirical norm and associated inner product

$$\|x_\theta - x_{\bar{\theta}}\|_n^2 := \frac{1}{n} \sum_{i=1}^n (x_{\theta_i} - x_{\bar{\theta}_i})^2 \qquad \langle x_\theta, x_{\bar{\theta}} \rangle_n := \frac{1}{n} \sum_{i=1}^n x_{\theta_i} x_{\bar{\theta}_i}$$

Then note that:

$$V_n(\theta) = \sum_{j=1}^M \theta_j \|x_j - x_\theta\|_n^2 \qquad Q(\theta) = (1-\nu) \|x_\theta - y\|_n^2 + \nu \sum_{j=1}^M \theta_j \|x_j - y\|_n^2 + \frac{\beta}{n} K(\theta)$$

From Taylor expansion on the Q-functional, we have for any $x_\theta, x_{\theta'} \in \mathbb{R}^n$:

$$Q(\theta) = Q(\theta') + \langle \nabla_x Q(\theta'), x_\theta - x_{\theta'} \rangle_n + (1-\nu) \|x_\theta - x_{\theta'}\|_n^2$$

Observe that for any $j \in \{1, \dots, M\}$:

$$Q(\theta^{(k)} + a_{k+1}(e_{J^{(k+1)}}) - \theta^{(k)}) \leq Q(\theta^{(k)} + a_{k+1}(e_j - \theta^{(k)}))$$

Multiplying by θ_j and summing over all θ we get that for any $\theta \in \Theta$:

$$\begin{aligned} \sum_{j=1}^M \theta_j Q(\theta^{(k)} + a_{k+1}(e_{J^{(k+1)}}) - \theta^{(k)}) &\leq \sum_{j=1}^M \theta_j Q(\theta^{(k)} + a_{k+1}(e_j - \theta^{(k)})) \\ \iff Q(\theta^{(k)} + a_{k+1}(e_{J^{(k+1)}}) - \theta^{(k)}) &\leq \sum_{j=1}^M \theta_j Q(\theta^{(k)} + a_{k+1}(e_j - \theta^{(k)})) \\ \implies Q(\theta^{(k+1)}) &\leq \sum_{j=1}^M \theta_j Q(\theta^{(k)} + a_{k+1}(e_j - \theta^{(k)})) \end{aligned}$$

Using the Taylor expansion formula with $\theta = \theta^{(k)} + a_{k+1}(e_j - \theta^{(k)})$ and $\theta' = \theta^{(k)}$ on the RHS:

$$\begin{aligned}
 Q(\theta^{(k+1)}) &\leq \sum_{j=1}^M \theta_j \left[Q(\theta^{(k+1)}) + [a_{k+1}(e_j - \theta^{(k)})]^T \nabla Q(\theta^{(k)}) + (1-\nu) \left\| \sum_{j=1}^M a_{k+1}(e_j - \theta^{(k)})x_j \right\|_n^2 \right] \\
 &\leq Q(\theta^{(k)}) + a_{k+1} \sum_{j=1}^M \theta_j \langle e_j, \nabla Q(\theta^{(k)}) \rangle_n - a_{k+1} \langle \theta^{(k)}, \nabla Q(\theta^{(k)}) \rangle_n + (1-\nu) a_{k+1}^2 \sum_{j=1}^M \theta_j \|x_j - x_{\theta^{(k)}}\|_n^2 \\
 &\leq Q(\theta^{(k)}) + a_{k+1} \langle \theta - \theta^{(k)}, \nabla Q(\theta^{(k)}) \rangle_n + (1-\nu) a_{k+1}^2 \sum_{j=1}^M \theta_j \|x_j - x_{\theta^{(k)}}\|_n^2 \\
 &\leq Q(\theta^{(k)}) + a_{k+1} [Q(\theta) - Q(\theta^{(k)})] - (1-\nu) a_{k+1} \|x_{\theta^{(k)}} - x_\theta\|_n^2 + (1-\nu) a_{k+1}^2 \sum_{j=1}^M \theta_j \|x_j - x_{\theta^{(k)}}\|_n^2 \\
 &\leq Q(\theta^{(k)}) + a_{k+1} [Q(\theta) - Q(\theta^{(k)})] - (1-\nu) a_{k+1} \|x_{\theta^{(k)}} - x_\theta\|_n^2 \\
 &\quad + (1-\nu) a_{k+1}^2 \left(\sum_{j=1}^M \theta_j \|x_j - x_\theta\|_n^2 + \|x_{\theta^{(k)}} - x_\theta\|_n^2 \right) \\
 &\leq Q(\theta^{(k)}) + a_{k+1} [Q(\theta) - Q(\theta^{(k)})] + (1-\nu) a_{k+1}^2 \sum_{j=1}^M \theta_j \|x_j - x_\theta\|_n^2
 \end{aligned}$$

$$Q(\theta^{(k+1)}) - Q(\theta) \leq (1 - a_{k+1})[Q(\theta^{(k)}) - Q(\theta)] + (1 - \nu) a_{k+1}^2 V(\theta)$$

For simplicity, denote $\delta_k = Q(\theta^{(k)}) - Q(\theta)$ and $B = (1 - \nu) \sum_{j=1}^M \theta_j \|x_j - x_\theta\|_n^2$

Using an induction on k, it is straightforward to prove the following:

$$\delta_k \leq \frac{4B}{k+3}$$

□

I.2 Excess Risk Guarantees for the k-step Greedy Q-aggregation Estimator

Theorem 32. Assume that for some function $\text{ERROR}(\hat{g})$:

$$P(\tilde{\ell}_{\theta, g_0}(Z) - \tilde{\ell}_{\theta', g_0}(Z)) - P(\tilde{\ell}_{\theta, \hat{g}}(Z) - \tilde{\ell}_{\theta', \hat{g}}(Z)) \leq \text{ERROR}(\hat{g}) \|f_\theta - f_{\theta'}\|^\gamma$$

for some $\gamma \in [0, 1]$ and square loss $\ell_{f, g}$. Assume that $\hat{Y}(\hat{g})$ and all functions f_j are absolutely uniformly bounded by U . Let θ^k be the k -step greedy approximate of $\hat{\theta}$. For any $k \geq 2$, $\nu \geq 1/2$ and $\mu \leq \frac{1}{60}(1 - \nu)$ and $\beta \leq \max\{\frac{64U^2(1-\nu)^2}{\mu}, \frac{16\nu U^2(\nu+8\mu U)}{\mu}\}$, w.p. $1 - \delta$:

$$R(\theta^k, g_0) \leq \min_{j=1}^M \left(R(e_j, g_0) + \frac{3\beta}{2n} \log(1/\pi_j) \right) + O\left(\frac{\log(1/\delta) \frac{1}{\mu} U^2}{n}\right) + \left(\frac{1}{\mu}\right)^{\frac{\gamma}{2-\gamma}} \text{ERROR}(\hat{g})^{\frac{2}{2-\gamma}}$$

Here, we prove an analogue of Lemma 18 for the K-step greedy approximate of $\hat{\theta}$.

Lemma 33. Let $\theta^* = \arg \min_{\theta \in \Theta} \tilde{R}(\theta, g_0) + C_V V(\theta) + C_K K(\theta)$ and let θ^k be the k -step greedy approximate of $\hat{\theta}$.

Define the centered and offset empirical processes:

$$\begin{aligned}
 Z_{n,1} &= (1 - \nu)(P - P_n)(\ell_{\theta^k, \hat{g}}(Z) - \ell_{\theta^*, \hat{g}}(Z)) - \mu \sum_j \theta_j^k \|f_j - f_{\theta^*}\|^2 - \frac{1}{s} K(\theta^k) \\
 Z_{n,2} &= \nu(P - P_n) \sum_j (\theta_j^k - \theta_j^*) \ell_{e_j, \hat{g}}(Z) - \mu(\theta^k)^\top H \theta^* - \frac{1}{s} K(\theta^k) & H_{ij} &= \|f_i - f_j\|^2 \\
 Z_{n,5} &= C_Q(V_n(\theta^*) - V(\theta^*)) - \mu V(\theta^*) - \frac{1}{s} K(\theta^*)
 \end{aligned}$$

where s is a positive number. Assume that:

$$\mathcal{E}(\hat{g}) := P(\tilde{\ell}_{\theta^k, g_0}(Z) - \tilde{\ell}_{\theta^*, g_0}(Z)) - P(\tilde{\ell}_{\theta^k, \hat{g}}(Z) - \tilde{\ell}_{\theta^*, \hat{g}}(Z)) \leq \text{ERROR}(\hat{g}) \|f_{\theta^k} - f_{\theta^*}\|$$

and $\ell_{f,g} = (\hat{Y})(g) - f$ where $\hat{Y}(\hat{g})$ and all functions f_j are absolutely uniformly bounded by U .

Choose $\beta \geq \frac{8n}{s}$, $C_K = \frac{3\beta}{2n}$, and $C_V < \min\{1 - \nu, \nu\}$. Furthermore, choose k such that $C_Q < C_V$. Then, assuming $\mu \leq \min\{\frac{1}{6}(1 - \nu - C_V), \frac{1}{4}(\nu - C_V), \frac{1}{2}(C_V - C_Q)\}$, we get:

$$R(\theta^k) \leq \min_j [R(e_j, g_0) + \frac{3\beta}{2n} \log(1/\pi_j)] + 2(Z_{n,1} + Z_{n,2} + Z_{n,5}) + \frac{1}{\mu} \text{ERROR}(\hat{g})^2$$

Proof. For k -step greedy approximate, Theorem 31 implies that:

$$\begin{aligned}
 P_n \tilde{\ell}_{\theta^k, \hat{g}}(Z) + \frac{\beta}{n} K(\theta^k) &\leq \min_{\theta \in \Theta} \{P_n \tilde{\ell}_{\theta, \hat{g}}(Z) + \frac{\beta}{n} K(\theta) + \frac{4(1 - \nu)}{k + 3} V_n(\theta)\} \\
 &\leq P_n \tilde{\ell}_{\theta^*, \hat{g}}(Z) + \frac{4(1 - \nu)}{k + 3} V_n(\theta^*) + \frac{\beta}{n} K(\theta^*) \tag{8}
 \end{aligned}$$

$$\begin{aligned}
 \tilde{R}(\theta^k, g_0) - \tilde{R}(\theta^*, g_0) &= P(\tilde{\ell}_{\theta^k, g_0}(Z) - \tilde{\ell}_{\theta^*, g_0}(Z)) \\
 &= P(\tilde{\ell}_{\theta^k, \hat{g}}(Z) - \tilde{\ell}_{\theta^*, \hat{g}}(Z)) + \mathcal{E}(\hat{g}) \\
 &\leq (P - P_n)(\tilde{\ell}_{\theta^k, \hat{g}}(Z) - \tilde{\ell}_{\theta^*, \hat{g}}(Z)) + \mathcal{E}(\hat{g}) - \frac{\beta}{n} (K(\theta^k) - K(\theta^*)) + C_Q V_n(\theta^*) \quad (\text{By Eqn 8}) \\
 &= \mu \sum_j \theta_j^k \|f_j - f_{\theta^*}\|^2 + \mu(\theta^k)^\top H \theta^* + Z_{n,1} + Z_{n,2} + Z_{n,5} + \mathcal{E}(\hat{g}) \\
 &\quad - \frac{\beta}{n} (K(\theta^k) - K(\theta^*)) + \frac{2}{s} K(\theta^k) + (C_Q + \mu) V(\theta^*) + \frac{1}{s} K(\theta^*) \\
 &\hspace{15em} (\text{definition of } Z_{n,1}, Z_{n,2}, Z_{n,5})
 \end{aligned}$$

By simple algebraic manipulations note that for any $\theta, \theta' \in \Theta$:

$$\begin{aligned}
 \sum_j \theta_j \|f_j - f_{\theta'}\|^2 &= V(\theta) + \|f_\theta - f_{\theta'}\|^2 \\
 \theta^\top H \theta' &= V(\theta) + V(\theta') + \|f_\theta - f_{\theta'}\|^2
 \end{aligned}$$

Letting $Z_n = Z_{n,1} + Z_{n,2} + Z_{n,5}$, we have:

$$\begin{aligned}
 \tilde{R}(\theta^k, g_0) - \tilde{R}(\theta^*, g_0) &\leq 2\mu V(\theta^k) + (2\mu + C_Q) V(\theta^*) + 2\mu \|f_{\theta^k} - f_{\theta^*}\|^2 + Z_n + \mathcal{E}(\hat{g}) \\
 &\quad - \frac{\beta}{n} (K(\theta^k) - K(\theta^*)) + \frac{2}{s} K(\theta^k) + \frac{1}{s} K(\theta^*) \\
 &\leq 2\mu V(\theta^k) + (2\mu + C_Q) V(\theta^*) + 2\mu \|f_{\theta^k} - f_{\theta^*}\|^2 + Z_n + \text{ERROR}(\hat{g}) \|f_{\theta^k} - f_{\theta^*}\|^\gamma \\
 &\quad - \frac{\beta}{n} (K(\theta^k) - K(\theta^*)) + \frac{2}{s} K(\theta^k) + \frac{1}{s} K(\theta^*)
 \end{aligned}$$

By Young's inequality, for any $a \geq 0$ and $b \geq 0$, we have $ab \leq \frac{a^p}{p} + \frac{b^q}{q}$, such that $\frac{1}{p} + \frac{1}{q} = 1$. Taking $a = \left(\frac{\gamma}{2\mu}\right)^{\gamma/2} \text{ERROR}(\hat{g})$ and $b = \left(\frac{2\mu}{\gamma}\right)^{\gamma/2} \|f_{\theta^k} - f_{\theta^*}\|^\gamma$ and $q = 2/\gamma$ and $p = \frac{1}{1-1/q} = \frac{1}{1-\gamma/2} = \frac{2}{2-\gamma}$

$$\text{ERROR}(\hat{g}) \|f_{\theta^k} - f_{\theta^*}\|^\gamma \leq \underbrace{\frac{2}{2-\gamma} \left(\frac{\gamma}{2\mu}\right)^{\frac{\gamma}{2-\gamma}} \text{ERROR}(\hat{g})^{\frac{2}{2-\gamma}} + \mu \|f_{\theta^k} - f_{\theta^*}\|^2}_{\mathcal{A}(\hat{g})}$$

Thus we can derive:

$$\begin{aligned} \tilde{R}(\theta^k, g_0) - \tilde{R}(\theta^*, g_0) &\leq 2\mu V(\theta^k) + (2\mu + C_Q)V(\theta^*) + 3\mu \|f_{\theta^k} - f_{\theta^*}\|^2 + Z_n + \mathcal{A}(\hat{g}) \\ &\quad - \frac{\beta}{n}(K(\theta^k) - K(\theta^*)) + \frac{2}{s}K(\theta^k) + \frac{1}{s}K(\theta^*) \end{aligned} \quad (9)$$

To obtain a lower bound:

$$\begin{aligned} \tilde{R}(\theta^k, g_0) - \tilde{R}(\theta^*, g_0) &= H(\theta^k, g_0) - H(\theta^*, g_0) + C_V(V(\theta^*) - V(\theta^k)) + C_K(K(\theta^*) - K(\theta^k)) \\ &\geq (1 - \nu - C_V) \|f_{\theta^k} - f_{\theta^*}\|^2 + C_V(V(\theta^*) - V(\theta^k)) + C_K(K(\theta^*) - K(\theta^k)) \end{aligned} \quad (\text{By Lemma 17})$$

Combining the latter upper and lower bound on $\tilde{R}(\theta^k, g_0) - \tilde{R}(\theta^*, g_0)$, we have:

$$\begin{aligned} (1 - \nu - C_V - 3\mu) \|f_{\theta^k} - f_{\theta^*}\|^2 &\leq (2\mu + C_V)V(\theta^k) + Z_n + \mathcal{A}(\hat{g}) + (2\mu + C_Q - C_V)V(\theta^*) \\ &\quad - \left(\frac{\beta}{n} - C_K\right) (K(\theta^k) - K(\theta^*)) + \frac{2}{s}K(\theta^k) + \frac{1}{s}K(\theta^*) \end{aligned}$$

Choosing:

$$(1 - \nu - C_V - 3\mu) \geq 3\mu \quad \Leftrightarrow \quad \frac{1}{6}(1 - \nu - C_V) \geq \mu$$

We get that:

$$\begin{aligned} 3\mu \|f_{\theta^k} - f_{\theta^*}\|^2 &\leq Z_n + \mathcal{A}(\hat{g}) + (2\mu + C_V)V(\theta^k) + (2\mu + C_Q - C_V)V(\theta^*) \\ &\quad - \left(\frac{\beta}{n} - C_K\right) (K(\theta^k) - K(\theta^*)) + \frac{2}{s}K(\theta^k) + \frac{1}{s}K(\theta^*) \end{aligned}$$

Plugging the bound on $\|f_{\theta^k} - f_{\theta^*}\|^2$ back into Equation (9) we have:

$$\begin{aligned} \tilde{R}(\theta^k, g_0) - \tilde{R}(\theta^*, g_0) &\leq (4\mu + C_V)V(\theta^k) + (4\mu + 2C_Q - C_V)V(\theta^*) + 2Z_n + 2\mathcal{A}(\hat{g}) \\ &\quad - \left(\frac{2\beta}{n} - C_K\right) (K(\theta^k) - K(\theta^*)) + \frac{4}{s}K(\theta^k) + \frac{2}{s}K(\theta^*) \end{aligned}$$

By Lemma 16, we have:

$$\begin{aligned} R(\theta^k, g_0) &\leq \tilde{R}(\theta^*, g_0) + (4\mu + 2C_Q - C_V)V(\theta^*) + (4\mu + C_V - \nu)V(\theta^k) + 2Z_n + 2\mathcal{A}(\hat{g}) \\ &\quad + \left(\frac{2\beta}{n} - C_K + \frac{2}{s}\right) K(\theta^*) + \left(\frac{4}{s} - \frac{2\beta}{n} + C_K\right) K(\theta^k) \\ &\leq \tilde{R}(\theta^*, g_0) + (4\mu + 2C_Q - C_V)V(\theta^*) + 2Z_n + 2\mathcal{A}(\hat{g}) \quad (\text{Since } \mu \leq \frac{1}{4}(\nu - C_V)) \\ &\quad + \left(\frac{2\beta}{n} - C_K + \frac{2}{s}\right) K(\theta^*) + \left(\frac{4}{s} - \frac{2\beta}{n} + C_K\right) K(\theta^k) \end{aligned}$$

Let $C_K = \frac{3\beta}{2n}$ Assuming:

$$\begin{aligned}\frac{1}{s} &\leq \frac{\beta}{8n} \\ C_V &\geq C_Q \\ \mu &\leq \frac{1}{2}(C_V - C_Q)\end{aligned}$$

We get that:

$$\begin{aligned}4\mu + 2C_Q - C_V &\leq C_V \\ \frac{2\beta}{n} - C_K + \frac{2}{s} &\leq C_K \\ \frac{4}{s} - \frac{2\beta}{n} + C_K &\leq 0\end{aligned}$$

Thus we get:

$$\begin{aligned}R(\theta^k, g_0) &\leq \tilde{R}(\theta^*, g_0) + C_V V(\theta^*) + C_K K(\theta^*) + 2Z_n + 2\mathcal{A}(\hat{g}) \\ &\leq \min_j [\tilde{R}(e_j, g_0) + C_V V(e_j) + C_K \log(1/\pi_j)] + 2Z_n + 2\mathcal{A}(\hat{g}) \\ &= \min_j [R(e_j, g_0) + \frac{3\beta}{2n} \log(1/\pi_j)] + 2Z_n + 2\mathcal{A}(\hat{g}) \\ &\quad \text{(Since } \tilde{R}(e_j, g_0) = R(e_j, g_0) \text{ and } V(e_j) = 0 \text{ for all } e_j\text{)}\end{aligned}$$

□

Lemma 34. Assuming that $f \in F$ are uniformly and absolutely and bounded in $[-U, U]$, for any positive $s \leq \frac{n\mu}{2C_Q^2 U^2(1+\mu)}$:

$$\Pr\left(Z_{n,5} \geq \frac{\log(1/\delta)}{s}\right) \leq \delta$$

Proof. It suffices to show that:

$$I := \mathbb{E}[\exp sZ_{n,5}] \leq 1$$

Note first that:

$$\begin{aligned}Z_{n,5} &= C_Q(V_n(\theta^*) - V(\theta^*)) - \mu V(\theta^*) - \frac{1}{s}K(\theta^*) \\ &\leq \max_j \left[C_Q(P_n - P)(f_j(Z) - f_{\theta^*}(Z))^2 - \mu P(f_j(Z) - f_{\theta^*}(Z))^2 + \frac{1}{s} \log(\pi_j) \right]\end{aligned}$$

Thus:

$$\begin{aligned}\mathbb{E}[\exp sZ_{n,5}] &\leq \mathbb{E} \left[\exp \left\{ s \max_j \left[C_Q(P_n - P)(f_j(Z) - f_{\theta^*}(Z))^2 - \mu P(f_j(Z) - f_{\theta^*}(Z))^2 + \frac{1}{s} \log(\pi_j) \right] \right\} \right] \\ &\leq \mathbb{E} \left[\max_j \exp \left\{ s \left(C_Q(P_n - P)(f_j(Z) - f_{\theta^*}(Z))^2 - \mu P(f_j(Z) - f_{\theta^*}(Z))^2 + \frac{1}{s} \log(\pi_j) \right) \right\} \right] \\ &\leq \sum_j \pi_j \mathbb{E} \left[\exp \left\{ s \left(C_Q(P_n - P)(f_j(Z) - f_{\theta^*}(Z))^2 - \mu P(f_j(Z) - f_{\theta^*}(Z))^2 \right) \right\} \right] \\ &\leq \max_j \mathbb{E} \left[\exp \left\{ s \left(C_Q(P_n - P)(f_j(Z) - f_{\theta^*}(Z))^2 - \mu P(f_j(Z) - f_{\theta^*}(Z))^2 \right) \right\} \right]\end{aligned}$$

Since $f_j(Z) \in [U, U]$, we have that $f_j(Z) - f_{\theta^*}(Z) \in [-2U, 2U]$. Thus:

$$(f_j(Z) - f_{\theta^*}(Z))^2 \geq \frac{(f_j(Z) - f_{\theta^*}(Z))^2}{4U^2} (f_j(Z) - f_{\theta^*}(Z))^2 = \frac{1}{4U^2} (f_j(Z) - f_{\theta^*}(Z))^4.$$

Thus:

$$I \leq \max_j \mathbb{E}_{Z_{1:n}} \left[\exp \left\{ s \left(C_Q (P_n - P) (f_j(Z) - f_{\theta^*}(Z))^2 - \frac{\mu}{4C_Q^2 U^2} PC_Q^2 (f_j(Z) - f_{\theta^*}(Z))^4 \right) \right\} \right]$$

It suffices to show that each term j is upper bounded by 1.

Applying Lemma 23 with:

$$\begin{aligned} V_i &= C_Q (f_j(Z_i) - f_{\theta^*}(Z_i))^2 \\ c &= 4C_Q^2 U^2 \\ \lambda &= s \\ c_0 &= \frac{\mu}{4C_Q^2 U^2} \end{aligned}$$

we have that as long as:

$$\lambda := \frac{s}{n} \leq \frac{2c_0}{1 + 2c_0 c} =: \frac{2 \cdot \mu / 4C_Q^2 U^2}{1 + \mu} = \frac{\mu}{2C_Q^2 U^2 (1 + \mu)}$$

Then:

$$\mathbb{E}_{Z_{1:n}} \left[\exp \left\{ s \left(C_Q (P_n - P) (f_j(Z) - f_{\theta^*}(Z))^2 - \frac{\mu}{4C_Q^2 U^2} PC_Q^2 (f_j(X) - f_{\theta^*}(X))^4 \right) \right\} \right] \leq 1$$

Thus we need that:

$$s \leq \frac{n\mu}{2C_Q^2 U^2 (1 + \mu)}$$

which is satisfied by our assumptions on s . □

Putting it All Together Choose $\nu \geq 1/2$, such that $1 - \nu \leq \nu$. Note that for any $k \geq 2$, we have $C_Q = \frac{4}{k+3}(1 - \nu) \leq \frac{4}{5}(1 - \nu)$. Choosing $C_V = \frac{5}{6}(1 - \nu) \leq \frac{5}{6}\nu$, then we have that for $k \geq 2$, we have $C_Q < C_V$ and we can take $\mu \leq \frac{1}{60}(1 - \nu)$ in Lemma 33. Combining all the above Lemmas, we have that w.p. $1 - \delta$, the k -step greedy Q-aggregation estimator θ_k achieves:

$$R(\theta_k) \leq \min_j \left(R(e_j) + \frac{3\beta}{2n} \log(1/\pi_j) \right) + O \left(\frac{\log(1/\delta) \frac{1}{\mu} U^2}{n} \right) + \frac{2}{\mu} \text{ERROR}(\hat{g})^2$$

J PROOF OF ORTHOGONAL LEARNING THEOREM

J.1 Fast Rates

Proof of Theorem 10. The proof follows the proof of Theorem 1 in Foster and Syrgkanis (2023) after replacing f_* by f , and we recover the original theorem by assuming $F' = F$.

By a second order Taylor expansion around f' , there exist $\bar{f} \in \text{Star}(F', f')$ such that:

$$\begin{aligned} (P\ell_{\hat{f}, \hat{g}} - P\ell_{f', \hat{g}}) - D_f P\ell(f', \hat{g})[\hat{f} - f'] &= \frac{1}{2} D_f^2 P\ell_{\bar{f}, \hat{g}}[\hat{f} - f', \hat{f} - f'] \\ &\geq \frac{\lambda}{2} \|\hat{f} - f'\|^2 - \frac{\kappa}{2} d(\hat{g}, g_0)^{\frac{4}{1+\tau}} \end{aligned} \quad (\text{By Assumption 4})$$

Rearranging, we get:

$$\frac{\lambda}{2} \|\hat{f} - f'\|^2 \leq (P\ell_{\hat{f}, \hat{g}} - P\ell_{f', \hat{g}}) - D_f P\ell(f', \hat{g})[\hat{f} - f'] + \frac{\kappa}{2} d(\hat{g}, g_0)^{\frac{4}{1+\tau}}$$

Applying a second-order Taylor expansion, we get:

$$\begin{aligned}
 -D_f P\ell(f', \hat{g})[\hat{f} - f'] &= -D_f P\ell(f', g_0)[\hat{f} - f'] - D_g D_f P\ell_{f', g_0}[\hat{f} - f', \hat{g} - g_0] \\
 &\quad - \frac{1}{2} D_g^2 D_f P\ell_{f, \hat{g}}[\hat{f} - f', \hat{g} - g_0, \hat{g} - g_0] \\
 &= -D_f P\ell(f', g_0)[\hat{f} - f'] - \frac{1}{2} D_g^2 D_f P\ell_{f, \hat{g}}[\hat{f} - f', \hat{g} - g_0, \hat{g} - g_0] \quad (\text{By Assumption 1}) \\
 &\leq -D_f P\ell(f', g_0)[\hat{f} - f'] + \beta_2 \|\hat{f} - f'\|^{1-r} d(\hat{g}, g_0)^2 \quad (\text{By Assumption 3(b)})
 \end{aligned}$$

Alternatively, we can get the above result directly by 5 instead of Assumptions 1 and 3.

Invoking Young's inequality and using $r \in [0, 1)$, we have that for any $\eta > 0$:

$$-D_f P\ell(f', \hat{g})[\hat{f} - f'] \leq -D_f P\ell(f', g_0)[\hat{f} - f'] + \frac{\beta_2 \eta}{4} \|\hat{f} - f'\|^2 + \frac{\beta_2}{2\eta^{(1-r)/(1+r)}} d(\hat{g}, g_0)^{\frac{4}{1+r}}$$

Choosing $\eta = \frac{\lambda}{\beta_2}$, we get that:

$$\|\hat{f} - f'\|^2 \leq -\frac{4}{\lambda} D_f P\ell(f', g_0)[\hat{f} - f'] + (P\ell_{\hat{f}, \hat{g}} - P\ell_{f', \hat{g}}) + O\left(d(\hat{g}, g_0)^{\frac{4}{1+r}}\right) \quad (10)$$

Taking Taylor expansion at f' , we have that there exist $\bar{f} \in \text{Star}(F', f')$ such that:

$$\begin{aligned}
 P\ell_{\hat{f}, g_0} - P\ell_{f', g_0} &= D_f P\ell_{f', g_0}[\hat{f} - f'] + \frac{1}{2} D_{\bar{f}}^2 P\ell_{\bar{f}, g_0}[\hat{f} - f', \hat{f} - f'] \\
 &\leq D_f P\ell_{f', g_0}[\hat{f} - f'] + \frac{\beta_1}{2} \|\hat{f} - f'\|^2 \quad (\text{By Assumption 3(a)}) \\
 &\leq -\left(\frac{2\beta_1}{\lambda} - 1\right) D_f P\ell_{f', g_0}[\hat{f} - f'] + O\left((P\ell_{\hat{f}, \hat{g}} - P\ell_{f', \hat{g}}) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)
 \end{aligned}$$

Since β_1 is the upper bound, we may assume $\frac{2\beta_1}{\lambda} > 1$. Thus, we get:

$$P\ell_{\hat{f}, g_0} - P\ell_{f', g_0} \leq O\left((P\ell_{\hat{f}, \hat{g}} - P\ell_{f', \hat{g}}) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Moreover, if

$$P\ell_{\hat{f}, \hat{g}}(Z) - P\ell_{f', \hat{g}}(Z) \leq \epsilon_n(\delta) \|\hat{f} - f'\| + \alpha_n(\delta)$$

Plugging this into Equation 10 yields:

$$\begin{aligned}
 \|\hat{f} - f'\|^2 &\leq -\frac{4}{\lambda} D_f P\ell(f', g_0)[\hat{f} - f'] + \epsilon_n(\delta) \|\hat{f} - f'\| + \alpha_n(\delta) + O\left(d(\hat{g}, g_0)^{\frac{4}{1+r}}\right) \\
 &\leq -\frac{4}{\lambda} D_f P\ell(f', g_0)[\hat{f} - f'] + \frac{1}{2} \|\hat{f} - f'\|^2 + O\left(\epsilon_n(\delta)^2 + \alpha_n(\delta) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)
 \end{aligned}$$

Thus, we get:

$$\|\hat{f} - f'\|^2 \leq -\frac{8}{\lambda} D_f P\ell(f', g_0)[\hat{f} - f'] + O\left(\epsilon_n(\delta)^2 + \alpha_n(\delta) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

Plugging this into the Taylor expansion at f' , we get:

$$\begin{aligned}
 P\ell_{\hat{f}, g_0} - P\ell_{f', g_0} &= D_f P\ell_{f', g_0}[\hat{f} - f'] + \frac{1}{2} D_{\bar{f}}^2 P\ell_{\bar{f}, g_0}[\hat{f} - f', \hat{f} - f'] \\
 &\leq D_f P\ell_{f', g_0}[\hat{f} - f'] + \frac{\beta_1}{2} \|\hat{f} - f'\|^2 \quad (\text{By Assumption 3(a)}) \\
 &\leq -\left(\frac{4\beta_1}{\lambda} - 1\right) D_f P\ell_{f', g_0}[\hat{f} - f'] + O\left(\epsilon_n(\delta)^2 + \alpha_n(\delta) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)
 \end{aligned}$$

Again, we may assume $\frac{4\beta_1}{\lambda} > 1$. Finally, we get:

$$P\ell_{\hat{f}, g_0} - P\ell_{f', g_0} \leq O\left(\epsilon_n(\delta)^2 + \alpha_n(\delta) + d(\hat{g}, g_0)^{\frac{4}{1+r}}\right)$$

□

J.2 Slow Rates

J.2.1 With Universal Orthogonality

Proof of Theorem 13. The proof follows the proof of Theorem 2 in Foster and Syrgkanis (2023), we reproduce it here for completeness.

$$P\ell_{\hat{f},g_0} - P\ell_{f',g_0} = (P\ell_{\hat{f},g_0} - P\ell_{\hat{f},\hat{g}}) + (P\ell_{f',\hat{g}} - P\ell_{f',g_0}) + (P\ell_{\hat{f},\hat{g}} - P\ell_{f',\hat{g}})$$

Taking a second order Taylor expansion yields that there exist $g, \bar{g} \in \text{Star}(G, g_0)$ such that:

$$\begin{aligned} (P\ell_{\hat{f},g_0} - P\ell_{\hat{f},\hat{g}}) + (P\ell_{f',\hat{g}} - P\ell_{f',g_0}) &= -D_g P\ell_{\hat{f},g_0}[\hat{g} - g_0] - \frac{1}{2}D_g^2 P\ell_{\hat{f},g_0}[\hat{g} - g - 0, \hat{g} - g_0] \\ &\quad + D_g P\ell_{f',g_0}[\hat{g} - g_0] + \frac{1}{2}D_g^2 P\ell_{f',g_0}[\hat{g} - g - 0, \hat{g} - g_0] \\ &\leq -D_g P\ell_{\hat{f},g_0}[\hat{g} - g_0] + D_g P\ell_{f',g_0}[\hat{g} - g_0] + \beta d(\hat{g}, g_0)^2 \end{aligned}$$

(By Assumption 7)

Next, we take another second-order Taylor expansion on the derivative terms:

$$\begin{aligned} D_g P\ell_{\hat{f},g_0}[\hat{g} - g_0] - D_g P\ell_{f',g_0}[\hat{g} - g_0] &= D_f D_g P\ell_{f',g_0}[\hat{g} - g_0, \hat{f} - f'] + \frac{1}{2}D_f^2 D_g P\ell_{f',g_0}[\hat{g} - g_0, \hat{f} - f', \hat{f} - f'] \\ &= \frac{1}{2}D_f^2 D_g P\ell_{f',g_0}[\hat{g} - g_0, \hat{f} - f', \hat{f} - f'] \end{aligned}$$

(By Assumption 6)

for some $\bar{f} \in \text{Star}(F, f')$. Furthermore, we can rewrite:

$$\begin{aligned} &D_f^2 D_g P\ell_{\bar{f},g_0}[\hat{g} - g_0, \hat{f} - f', \hat{f} - f'] \\ &= \lim_{t \rightarrow 0} \frac{D_f D_g P\ell_{\bar{f}+t(\hat{f}-f'),g_0}[\hat{g} - g_0, \hat{f} - f'] - D_f D_g P\ell_{\bar{f},g_0}[\hat{g} - g_0, \hat{f} - f']}{t} \end{aligned}$$

Since $\bar{f} + t(\hat{f} - f') \in \text{Star}(F, f') + \text{Star}(F - f', 0)$ for all $t \in [0, 1]$, Assumption 6 implies that both terms in the numerator are zero. Thus, we get:

$$D_g P\ell_{\hat{f},g_0}[\hat{g} - g_0] = D_g P\ell_{f',g_0}[\hat{g} - g_0]$$

Combining, we arrive at the desired result:

$$P\ell_{\hat{f},g_0}(Z) - P\ell_{f',g_0}(Z) \leq O\left(P\ell_{\hat{f},\hat{g}}(Z) - P\ell_{f',\hat{g}}(Z) + d(g, g_0)^2\right)$$

□

J.2.2 Without Universal Orthogonality

Proof of Theorem 13.

$$\begin{aligned} P\ell_{\hat{f},g_0} - P\ell_{f',g_0} &= (P\ell_{\hat{f},g_0} - P\ell_{f',g_0}) - (P\ell_{\hat{f},\hat{g}} - P\ell_{f',\hat{g}}) + (P\ell_{\hat{f},\hat{g}} - P\ell_{f',\hat{g}}) \\ &= D_f P\ell(\bar{f}, g_0)[\hat{f} - f'] - D_f P\ell(\bar{f}, \hat{g})[\hat{f} - f'] + P\ell_{\hat{f},\hat{g}} - P\ell_{f',\hat{g}} \\ &\quad \text{(By first order Taylor expansion)} \\ &\leq \beta_4 d(\hat{g}, g_0)^2 + (P\ell_{\hat{f},\hat{g}} - P\ell_{f',\hat{g}}) \quad \text{(By Assumption 8)} \\ &= O\left(P\ell_{\hat{f},\hat{g}} - P\ell_{f',\hat{g}} + d(\hat{g}, g_0)^2\right) \end{aligned}$$

□