

---

# Kernel Debiased Plug-in Estimation: Simultaneous, Automated Debiasing without Influence Functions for Many Target Parameters

---

Brian Cho<sup>1</sup> Yaroslav Mukhin<sup>2</sup> Kyra Gan<sup>1</sup> Ivana Malenica<sup>3</sup>

## Abstract

When estimating target parameters in nonparametric models with nuisance parameters, substituting the unknown nuisances with nonparametric estimators can introduce “plug-in bias.” Traditional methods addressing this suboptimal bias-variance trade-off rely on the *influence function* (IF) of the target parameter. When estimating multiple target parameters, these methods require debiasing the nuisance parameter multiple times using the corresponding IFs, which poses analytical and computational challenges. In this work, we leverage the *targeted maximum likelihood estimation* (TMLE) framework to propose a novel method named *kernel debiased plug-in estimation* (KDPE). KDPE refines an initial estimate through regularized likelihood maximization steps, employing a nonparametric model based on *reproducing kernel Hilbert spaces*. We show that KDPE: (i) simultaneously debiases *all* pathwise differentiable target parameters that satisfy our regularity conditions, (ii) does not require the IF for implementation, and (iii) remains computationally tractable. We numerically illustrate the use of KDPE and validate our theoretical results.

## 1. Introduction

Estimating a target parameter  $\psi(P)$  from  $n$  independent samples of an unknown distribution  $P$  is a fundamental and rapidly evolving statistical field. Driven by data-adaptive machine learning methods and efficiency theory, modern approaches to estimation can achieve optimal performance under minimal assumptions on the true data-generating process (DGP). These statistical learning advancements enable researchers to obtain strong theoretical guarantees and empirical performance while avoiding unrealistic assumptions.

---

<sup>1</sup>Department of ORIE, Cornell Tech, NY, USA <sup>2</sup>Massachusetts Institute of Technology, Cambridge, MA, USA <sup>3</sup>Department of Statistics, Harvard University, Cambridge, MA, USA. Correspondence to: Brian Cho <bmc233@cornell.edu>.

Modern estimation methods rely on the plug-in principle (van der Vaart, 2000), which substitutes unknown parameters of the underlying data-generating process with estimated empirical counterparts (e.g., the sample mean is the average of observations over the empirical distribution  $\mathbb{P}_n$  defined by the dataset). Flexible, machine learning (ML) estimation methods have further exploited the plug-in approach. Meta-learners in the causal inference literature use random forests and neural networks to estimate empirical conditional density/regression functions for the downstream task of estimating target parameters such as potential outcome means and treatment effects (Künzel et al., 2019). The use of highly adaptive, complex ML algorithms, however, induces *plug-in bias* (first-order bias) that impacts the downstream estimate.

Estimators such as double machine learning (DML) (Chernozhukov et al., 2017), targeted maximum likelihood estimation (van der Laan and Rubin, 2006), and the one-step correction (Bickel et al., 1998) remove plug-in bias and achieve efficiency (i.e., lowest possible asymptotic variance) within the class of *regular and asymptotically linear* (RAL) estimators. However, these methods typically require knowledge of the influence function or the efficient influence function (EIF)<sup>1</sup> of the target parameter. Deriving and computing these functions analytically is often challenging for practitioners (Carone et al., 2018; Jordan et al., 2022; Kennedy, 2022) due to their dependence on both the statistical functional of interest and the statistical model (e.g., assumptions on the DGP). For example, even seemingly innocuous target parameters, such as the average density value, require specialized knowledge to study with conventional techniques (Carone et al., 2018). Further, the knowledge of the IF/EIF is specific to a single target parameter under a particular statistical model. When estimating multiple target parameters, existing methods require the derivation of the IF/EIF for *each target parameter*, hindering the adoption of these techniques in practice (Hines et al., 2022).

**Contributions** We propose a generic approach for constructing efficient RAL estimators that combines the TMLE

---

<sup>1</sup>The notation of EIF arises when we are in the semi-parametric setting where the directions in which we can perturb the data generating distribution are restricted. Otherwise, IF is the EIF.

(van der Laan and Gruber, 2016) framework with a novel application of the reproducing kernel Hilbert spaces (RKHSs). Leveraging the universal approximation property of RKHSs, we construct a debiased distribution (or nuisance) estimate  $P_n^\infty$  that has *negligible plug-in bias* and *attains efficiency* under similar assumptions to existing approaches. Contrary to popular approaches, our kernel debiased plug-in approach

1. offers an *automated framework* that does not require the computation of (efficient) influence function  $\phi_P$ , debiased/orthogonalized estimating equations, or efficiency bounds from the user, and
2. the same nuisance/distribution estimate  $P_n^\infty$  can be used as a plug-in estimate for *all* pathwise differentiable target parameters  $\psi(P)$  that satisfy some standard regularity conditions (provided in Section 3).

The technical contribution of our work is to identify the set of sufficient conditions to eliminate the plug-in bias under our proposed framework. In particular, Assumption 5 is our novel regularity assumption on the model  $\mathcal{M}$  and functional  $\psi$  needed to control the plug-in bias term of the proposed KDPE estimator. We numerically validate our results and illustrate that our method performs as well as state-of-the-art modern causal inference methods, which explicitly use the functional form of the IF.

The paper is organized as follows. In Section 2, we describe the problem setup and introduce key concepts related to TMLE and RKHSs. To ease presentation, we focus on the most generic case where 1) the estimation problem is non-parametric (i.e., no assumption on the DGP), 2) the nuisance is considered to be the entire distribution  $P$ , and 3) the target parameter  $\psi$  is a mapping from the probability distribution  $P$  to real numbers. We extend our procedure to semiparametric models and other nuisances in Appendix C.2. In Section 3, we propose *kernel debiased plug-in estimator*, and characterize the assumptions needed for KDPE to be asymptotically linear and efficient. In Section 4, we provide concrete empirical examples of KDPE for target parameters such as the average treatment effect, odds ratio, and relative risk, in both single-time period and longitudinal settings.

### 1.1. Related Works

Our work is related to two types of approaches that can debias a target parameter: 1) IF-based, and 2) IF-free computational approaches. When compared with these two approaches, *KDPE simultaneously debiases multiple target parameters* and does not require the derivation of IF. In addition, KDPE relates to undersmoothed *highly adaptive lasso* (HAL-) MLE, which can be used to simultaneously debias multiple parameters. When compared with HAL-MLE, our method is computationally more efficient.

The main techniques for constructing IF-based efficient RAL estimators include estimating equations, DML, one-

step correction, and TMLE. A general estimating equation approach involves solving for the target parameter by setting the score equations to zero (Robins, 1986; Newey, 1990; Bickel et al., 1998). Double machine learning characterizes the estimand as the solution to a population score equation (Chernozhukov et al., 2017). One-step correction method adds the empirical average of the IF to an initial estimator (Bickel et al., 1998). We defer the discussion on TMLE to Section 2.3.

Current IF-free methods for efficient estimation include a) approximating IF through finite-differencing (Carone et al., 2018; Jordan et al., 2022) or Monte Carlo (Agrawal et al., 2024), and b) AutoDML (Chernozhukov et al., 2022). In contrast to methods in a), KDPE does not attempt to approximate the IF of a target parameter in any way. AutoDML, on the other hand, relies on an orthogonal reparameterization of the problem to achieve debiasing. It automates the estimation of the additional nuisance parameter by exploiting the structure of the estimating equations in the particular setting of Chernozhukov et al. (2022). KDPE, in contrast, directly considers the plug-in bias and eliminates this term from the final estimate.

The HAL-MLE (van der Laan et al., 2022) solves the score equation for *all* cadlag functions with bounded L1 norm (which is assumed to include the desired EIF). Like HAL, KDPE solves a rich set of scores, rather than a single targeted score, to approximately solve the score equation at a  $\sqrt{n}$ -rate. HAL, however, is highly computationally inefficient for large models due to its basis functions growing exponentially with the number of covariates  $d$  and polynomially with the sample size  $n$ . KDPE uses only  $n$  basis functions, improving computational tractability to solve all scores in a *universal* RKHS.

## 2. Problem Setup and Preliminaries

Let  $O_1, \dots, O_n$  denote iid observations drawn from a probability distribution  $P^*$  on the sample space  $\mathcal{O}$ . The distribution  $P^*$  is *unknown* but is assumed to belong to a *nonparametric*<sup>2</sup> collection  $\mathcal{M}$ , which consists of distributions on  $\mathcal{O}$  dominated by a common measure. Let  $\psi : \mathcal{M} \rightarrow \mathbb{R}^m$  be a functional of the model  $\mathcal{M}$ , also referred to as the *target parameter*. Our goal is to “efficiently” estimate  $\psi(P^*)$  on the basis of the  $n$  observations using a plug-in estimator (i.e., by estimating  $P^*$  and evaluating  $\psi$  on this estimate) *without using influence functions*. Since our method produces a single debiased estimate  $P_n^\infty$  that is *independent* of target parameter  $\psi$ , we assume that  $m = 1$  wlog.

**Notation** The density of  $P$  is denoted by the lowercase  $p$ , which we use interchangeably to refer to the probability

<sup>2</sup>We provide a simple semiparametric extension, where we restrict the data perturbation directions, in Appendix C.2.

measure. We use  $Pf$  and  $P[f]$  to indicate the expectation  $\mathbb{E}_P[f] = \int f dP$  and  $\mathbb{P}_n$  to denote the empirical measure  $\mathbb{P}_n f := \frac{1}{n} \sum_{i=1}^n f(O_i)$ . Let  $L_0^2(P)$  be the set of all square-integrable functions with respect to probability measure  $P$ , i.e.,  $f : \mathcal{O} \rightarrow \mathbb{R}$  with  $Pf = 0$  and  $Pf^2 < \infty$ . A complete notation table is provided in Appendix A.

To formalize the notation of asymptotic efficiency, we first introduce RAL estimators in Section 2.1. In Section 2.2, we introduce *plug-in bias* and provide the assumptions and condition required for plug-in estimators to be *asymptotically linear and efficient*. We introduce TMLE in Section 2.3 since KDPE shares the same construction framework as TMLE, with one main modification. Finally, we introduce the necessary concepts related to RKHSs in Section 2.4, as KDPE employs a RKHS-based model.

## 2.1. Regular and Asymptotic Linear Estimators

*Asymptotic linearity* of an estimator leads to a tractable limiting distribution, resulting in asymptotically valid inference. It corresponds to the ability to approximate the difference between an estimate and the true value of the target parameter as an average of i.i.d. random variables.

**Definition 1** (Asymptotic linearity). *An estimator  $\hat{\psi}_n$  is asymptotically linear if  $\hat{\psi}_n - \psi(P^*) = \mathbb{P}_n \phi_{P^*} + o_{P^*}(1/\sqrt{n})$ , where  $\phi_{P^*} : \mathcal{O} \rightarrow \mathbb{R}^m$  is the corresponding influence function for estimator  $\hat{\psi}_n$ .*

On the other hand, a *regular* estimator attains the same limiting distribution even under perturbations (on the order of  $\sqrt{n}$ ) of the true data-generating distribution, enforcing robustness to distributional shifts. Since the most efficient (i.e., smallest sampling variance) regular estimator is guaranteed to be asymptotically linear (van der Vaart, 1991), we restrict our attention to the class of RAL estimators.<sup>3</sup>

## 2.2. Plug-in Bias

To connect plug-in estimators (i.e.,  $\hat{\psi}_n \equiv \psi(\hat{P})$ , where  $\hat{P}$  is our distributional estimate with  $n$  samples) and RAL estimators, we consider target parameters that satisfy *pathwise differentiability* (Koshevnik and Levit, 1976). This means that given a smoothly parametrized one-dimensional submodel  $\{P_\epsilon\}_{0 < \epsilon < 0^+}$  of  $\mathcal{M}$ , the map  $\epsilon \mapsto \psi(P_\epsilon)$  is differentiable in the ordinary sense,<sup>4</sup> and the derivative has a Riesz representation discussed below. Pathwise differentiability of a target parameter admits *regular* estimators that (i) converge at the parametric  $\sqrt{n}$ -rate and (ii) whose asymptotic normal distributions vary smoothly with  $P$ . To

<sup>3</sup>Our nonparametric assumption on the model class  $\mathcal{M}$  implies that all RAL estimators have the IF.

<sup>4</sup>For a concrete definition of the one-dimensional sub-model  $\{P_\epsilon\}_{0 < \epsilon < 0^+}$ , we refer readers to Equation (1), which gives an explicit submodel  $p_{\epsilon,h}$ . There exist many choices for this submodel; we opt for the linear model  $p_{\epsilon,h}$  here for simplicity.

see this, we first introduce the tangent space.

**Scores** Let  $h = \frac{d}{d\epsilon}|_{\epsilon=0} \log p_\epsilon$  be the *score* of the path at  $P_0$  (i.e.,  $P_{\epsilon=0}$ ). Then, the smoothness requirement on the path implies that the expected value of the score function under distribution  $P_0$  equals 0, i.e.,  $P_0 h = 0$ , and  $P_0 h^2 < \infty$  (van der Vaart, 2000). The collection of scores of all smooth paths through  $P_0$  forms a linear space, which we refer to as the *tangent space* of  $\mathcal{M}$  at  $P_0$ . Assuming that  $\mathcal{M}$  is nonparametric, for any  $P \in \mathcal{M}$ , its tangent space equals  $L_0^2(P)$ .<sup>5</sup> Given any distribution  $P \in \mathcal{M}$  and a score  $h$ , we work with the concrete path<sup>6</sup>

$$p_{\epsilon,h} := p(\epsilon|h) := (1 + \epsilon h)p. \quad (1)$$

**Influence function** While the IF defines the limiting sampling variance of a RAL estimator in Definition 1, it has additional, related interpretations: the IF is the Riesz representer of the derivative functional. Pathwise differentiability of  $\psi$  at  $P$  implies that there exists a continuous linear functional<sup>7</sup>  $D_P \psi : L_0^2(P) \rightarrow \mathbb{R}$  such that for any smooth path with a score of  $h$ , the derivative of the functional  $\psi(\cdot)$  at distribution  $P$ ,  $\frac{d}{d\epsilon}|_{\epsilon=0} \psi(P_{\epsilon,h})$ , is equivalent to evaluating  $D_P \psi[\cdot]$  at score  $h$ , i.e.,  $\frac{d}{d\epsilon}|_{\epsilon=0} \psi(P_{\epsilon,h}) = D_P \psi[h]$ . The *influence function*  $\phi_P \in L_0^2(P)$  of parameter  $\psi$  is the Riesz representer of its derivative  $D_P \psi$  for the  $L_0^2(P)$  inner product and is characterized by the property that  $D_P \psi[h] = \langle h, \phi_P \rangle_{L_0^2(P)}$  for every  $h \in L_0^2(P)$ .

Next, Lemma 1 (Bickel et al., 1998) decomposes the estimation error,  $\psi(\hat{P}) - \psi(P^*)$ , by considering the following concrete path between  $\hat{P}$  and  $P^*$ :  $\hat{p}(\epsilon|h) = (1 - \epsilon)\hat{p} + \epsilon P^*$ . We note that under this path, the direction  $h$  becomes  $h = P^*/\hat{p} - 1$ . By leveraging the Riesz Representation property, we obtain the error decomposition shown below:

**Lemma 1** (Bickel et al. 1998). *Let  $\psi : \mathcal{M} \rightarrow \mathbb{R}$  be pathwise differentiable and  $\hat{P}, P^* \in \mathcal{M}$ . Assume that  $\hat{P}$  satisfies: (1)  $P^*$  is absolutely continuous with respect to  $\hat{P}$ , and (2) Radon–Nikodym derivative  $dP^*/d\hat{P}$  is square integrable under  $\hat{P}$ . Then the following von Mises expansion holds:*

$$\begin{aligned} \psi(\hat{P}) - \psi(P^*) &= \mathbb{P}_n \phi_{P^*} - \mathbb{P}_n \phi_{\hat{P}} \\ &\quad + (\mathbb{P}_n - P^*)[\phi_{\hat{P}} - \phi_{P^*}] + R_2(\hat{P}, P^*) \end{aligned} \quad (2)$$

where  $R_2(\hat{P}, P^*)$  is the second-order (quadratic) remainder in the difference between  $\hat{P}$  and  $P^*$ .

<sup>5</sup>Intuitively, the tangent space at distribution  $P$  contains all directions through which we can move from the current distribution  $P$  to another distribution in model class  $\mathcal{M}$ .

<sup>6</sup>Intuitively, the set of  $\{P_{\epsilon,h}\}_{\epsilon,h}$  can be thought as a reparameterization of the model class  $\mathcal{M}$  centered around a local distribution  $P$ . We refer to Appendix B.1 for further questions regarding the validity of our linear submodel.

<sup>7</sup>Linear with respect to  $h$ .

The expansion in Lemma 1 closely resembles Definition 1. It motivates the following common assumptions that control the behavior of the empirical process term  $(\mathbb{P}_n - P^*)[\phi_{\hat{P}} - \phi_{P^*}]$  and the second-order reminder  $R_2(\hat{P}, P^*)$  (van der Laan and Rubin, 2006).

**Assumption 1.** Assume that (i)  $\|\phi_{\hat{P}} - \phi_{P^*}\|_{L^2(P^*)} = o_{P^*}(1)$  and that (ii) there exists an event  $\Omega$  with  $P^*(\Omega) = 1$  such that the set of functions  $\{\phi_{\hat{P}_n(w)}; w \in \Omega, n \in \mathbb{N}\}$  is  $P^*$ -Donsker. Then the empirical process term satisfies:  $(\mathbb{P}_n - P^*)[\phi_{\hat{P}} - \phi_{P^*}] = o_{P^*}(1/\sqrt{n})$ .

**Assumption 2.** Assume that  $\hat{P}$  converges sufficiently fast and  $\psi$  is sufficiently regular so that the second-order remainder term  $R_2(\hat{P}, P^*)$  is  $o_{P^*}(1/\sqrt{n})$ .

We note that Assumptions 1 and 2 are standard in the TMLE literature (van der Laan and Rubin, 2006). The goal of this work is to construct a plug-in estimator  $\psi(\hat{P})$  that is asymptotically linear and converges at the parametric (efficient)  $\sqrt{n}$ -rate to the Gaussian asymptotic distribution  $N(0, P^*[\phi_{P^*}^2])$ . Since  $\mathbb{P}_n \phi_{P^*}$  governs the asymptotic distribution of  $\psi(\hat{P})$ , Assumptions 1 and 2 leave us with a single term that must converge at an  $o_{P^*}(1/\sqrt{n})$  rate:

$$\psi(\hat{P}) - \psi(P^*) = \mathbb{P}_n \phi_{P^*} - \mathbb{P}_n \phi_{\hat{P}} + o_{P^*}(1/\sqrt{n}).$$

We denote this remaining term  $\mathbb{P}_n \phi_{\hat{P}}$  as the *plug-in bias*. In contrast, a naive plug-in estimator may have a first-order/plug-in bias term that dominates the  $\sqrt{n}$ -asymptotics.

### 2.3. TMLE Preliminaries

We briefly recap TMLE (van der Laan and Rubin, 2006), as its construction closely resembles that of our estimator. TMLE is a plug-in estimator that satisfies Definition 1 and is constructed in two steps. First, one obtains a preliminary estimate  $P_n^0$  of  $P^*$ , which is typically consistent but not efficient (slower than  $\sqrt{n}$ -rate). The TMLE solution to debias  $P_n^0$  for the parameter  $\psi$  is to perturb  $P_n^0$  in a way that (i) increases the sample likelihood of the distribution estimate, and (ii) sets the plug-in bias  $\mathbb{P}_n \phi_{\hat{P}} = 0$ , maintaining consistency and mitigating first-order bias. van der Laan and Rubin (2006) consider the parametric model in (1) with the score equal to the influence function  $\phi_{P_n^0}$  and update the estimate to be the MLE  $p_n^1 = (1 + \epsilon_n \phi_{P_n^0}) P_n^0$ , where

$$\epsilon_n := \arg \max_{\epsilon: p_n^0(\epsilon) \in \mathcal{M}} \sum_{i=1}^n \log(1 + \epsilon \phi_{P_n^0}(O_i)) P_n^0(O_i). \quad (3)$$

Assuming that we converge to an interior point such that first-order conditions hold, the updated estimate  $p_n^1$  solves the following:

$$0 = \mathbb{P}_n \frac{d}{d\epsilon} p_n^0(\epsilon) = \frac{1}{n} \sum_{i=1}^n \frac{\phi_{P_n^0}(O_i)}{1 + \epsilon \phi_{P_n^0}(O_i)} = \mathbb{P}_n \left[ \frac{\phi_{P_n^0}}{1 + \epsilon \phi_{P_n^0}} \right]. \quad (4)$$

By iterating the TMLE step, we get a sequence of updated estimates  $\{p_n^\ell\}_{\ell=0}^\infty$ . Assuming that  $p_n^\ell$  converges as  $\ell \rightarrow \infty$  sufficiently regularly,<sup>8</sup> the limit  $p_n^\infty$  is a *fixed point* of (3), the corresponding argmax in (3) is  $\epsilon_n = 0$ , and (4) becomes the plug-in bias term:  $\mathbb{P}_n \phi_{P_n^\infty} = 0$ . This property of  $P_n^\infty$  achieves asymptotic linearity and efficiency for the plug-in estimator  $\psi(P_n^\infty)$  under Assumptions 1 and 2.

We construct KDPE following the TMLE framework with one modification. Instead of reducing  $\mathcal{M}$  into a one-dimensional submodel by perturbing  $P_n^0$  in the direction of  $\phi_{P_n^0}$ , we construct RKHS-based submodels. These submodels are of infinite dimensions and are independent of both  $\phi_P$  and  $\psi$ . Next, we introduce RKHS properties.

### 2.4. RKHS Preliminaries

Let  $\mathcal{X}$  be a non-empty set. A function  $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  is called positive-definite (PD) if for any finite sequence of inputs  $\mathbf{x} = [x_1, \dots, x_n]^\top \in \mathcal{X}^n$ , the matrix  $K_{\mathbf{x}} := [K(x_i, x_j)]_{ij}$  is symmetric and positive-definite (Micchelli et al., 2006). The *kernel function* at  $x$  is the univariate map  $y \mapsto k_x(y) := K(x, y)$ ; it is common to overload the notation with vector-valued maps  $k_{\mathbf{x}} := [k_{x_1}, \dots, k_{x_n}]^\top$ ; letting  $\mathbf{y} = [y_1, \dots, y_m]^\top$  be a column vector, we define the *kernel matrix*  $K_{\mathbf{xy}} := [K(x_i, y_j)]_{ij} \in \mathbb{R}^{n \times m}$  and  $K_{\mathbf{xx}} := K_{\mathbf{xx}}$ .

The *reproducing kernel Hilbert space* associated with kernel  $K$ , denoted  $\mathcal{H}_K$ , is a unique Hilbert space of functions on  $\mathcal{X}$  that satisfies the following properties: (i) for any  $x \in \mathcal{X}$ , the function  $k_x \in \mathcal{H}_K$ , and (ii)  $\langle f, k_x \rangle_{\mathcal{H}} = f(x)$  for all  $f \in \mathcal{H}_K$ . RKHSs have two key features that facilitate our results: (i) suitable choices of RKHSs, designated as universal kernels, are sufficiently rich to approximate any influence function, and (ii) optimization of sample-based objectives (e.g., likelihood) over the RKHS-based models can be efficiently carried out using a representer-type theorem, allowing for a tractable solution.

Denoting  $\mathcal{C}_0(\mathcal{X})$  as the space of continuous functions that vanish at infinity with the uniform norm, we say that a kernel  $K$  is  $\mathcal{C}_0$  provided that  $\mathcal{H}_K$  is a subspace of  $\mathcal{C}_0(\mathcal{X})$ . We say that a  $\mathcal{C}_0$ -kernel  $K$  is *universal* (Carmeli et al., 2010) if it satisfies the following property:

**Property 1** (Universal kernels). *Let  $\mathcal{X}$  be a closed subset of  $\mathbb{R}^n$ , and let  $K$  be a PD kernel such that the map  $x \mapsto K(x, x)$  is bounded and  $k_x \in \mathcal{C}_0(\mathcal{X})$  for all  $x \in \mathcal{X}$ . Then  $\mathcal{H}_K \subset \mathcal{C}_0(\mathcal{X}) \subset L^p(P) \subset L^q(P)$  for all  $1 \leq p < q \leq \infty$  and every probability measure  $P$ . Furthermore, the kernel  $K$  is universal if and only if the space  $\mathcal{H}_K$  is dense in  $L^p(\mathcal{X}, P)$  for at least one (and consequently for all)  $1 \leq p < \infty$  and each probability measure  $P$  on  $\mathcal{X}$ .*

<sup>8</sup>We have omitted details from the main exposition of TMLE to provide a brief overview of the method. We refer to the last paragraph on Page 11, (van der Laan and Rubin, 2006), for the full characterization.

Property 1 implies that universal kernels contain sequences that can approximate the influence function  $\phi_P$  arbitrarily well with respect to the  $L^2(P)$  norm for any  $P$ . Universal kernels include the *Gaussian kernel*  $K(x, y) = \exp(-\|x - y\|_2^2)$ , which is the primary example in our work. Our debiasing method relies on constructing local (to  $P$ ) submodels of  $\mathcal{M}$  indexed by a set of scores at  $P$ . Since scores are  $P$ -mean-zero, given any PD kernel  $K$  on  $\mathcal{X}$  that is bounded,  $C_0$ , and universal, we transform  $K$  into a mean-zero kernel with respect to  $P$  (proof in Appendix B.2):

**Proposition 1** (Mean-zero RKHS). *Let  $\mathcal{H}_P := \{h \in \mathcal{H}_K ; Ph = 0\}$  be the subspace of  $\mathcal{H}_K$  containing all  $P$ -mean-zero functions. Then 1)  $\mathcal{H}_P = \mathcal{H}_{K_P}$  is closed in  $\mathcal{H}_K$  and is also an RKHS, with reproducing kernel  $K_P(x, y) := K(x, y) - \frac{\int_{\mathcal{X}} K(x, s) dP(s) \int_{\mathcal{X}} K(y, s) dP(s)}{\int_{\mathcal{X} \times \mathcal{X}} K(s, t) dP(s) dP(t)}$ , and 2)  $\mathcal{H}_P$  is dense in  $L_0^2(P)$ .*

### 3. Debiasing with RKHS

We propose a novel distributional estimator  $P_n^\infty$ , which has the following properties: (i) it takes two user inputs: a pre-estimate  $P_n^0$  and an RKHS  $\mathcal{H}_K$  (equivalently, the kernel  $K$ ) on the sample space  $\mathcal{O}$ ; (ii) it provides a debiased and efficient plug-in  $\psi(P_n^\infty)$  for estimating any parameter  $\psi$  under the stated regularity conditions; (iii) it solves the influence curve estimating equation asymptotically, thereby eliminating the plug-in bias of  $P_n^0$  but unlike TMLE this holds simultaneously for all pathwise differentiable parameters; (iv) it does not require the influence function  $\phi_P$  to be implemented and does not depend on any  $\psi$ . We now describe each step of this estimator in detail.

**Local fluctuations of pre-estimate** Let  $\mathcal{H}_K$  be an RKHS with a bounded kernel  $K$  on the sample space  $\mathcal{O}$  and  $\mathcal{H}_P$  be the subspace of scores at  $P$  within  $\mathcal{H}$ . For each  $P \in \mathcal{M}$ ,

$$\begin{aligned} \widetilde{\mathcal{M}}_P &:= \{p_h := p(h) := (1 + h)p; \\ &h \in \mathcal{H}_P \text{ such that } (1 + h)p > 0\} \cap \mathcal{M}. \end{aligned} \quad (5)$$

This submodel<sup>9</sup> allows  $P$  to be perturbed along any direction  $h \in \mathcal{H}_P$  that results in a valid distribution in the model  $\mathcal{M}$ . These directions  $h$  must have zero mean under  $P$ , requiring modifying the kernel  $K$  as described in Proposition 1.

**Regularized MLE** To choose the debiasing update from  $\widetilde{\mathcal{M}}_P$ , we follow TMLE (Section 2.3) and maximize the empirical likelihood. However, since our fluctuations are indexed by the infinite-dimensional set  $\mathcal{H}_P$ , the MLE problem is generally not well-posed (Fukumizu, 2009) and must be regularized. This is achieved with an RKHS-norm penalty, leading to the following optimization problem as a mecha-

<sup>9</sup>We provide further details regarding the RKHS-constrained model  $\widetilde{\mathcal{M}}_P$  and its interpretation in Appendix B.1.

nism to choose the update for an estimate  $P$ :

$$\arg \min_{h \in \mathcal{H}_P : p(h) \in \widetilde{\mathcal{M}}_P} \mathbb{P}_n [-\log(1 + h)p] + \lambda \|h\|_{\mathcal{H}_P}^2, \quad (6)$$

where the regularization parameter  $\lambda \geq 0$  regularizes the complexity of the perturbation  $h$  to  $P$ . Note that (6) is an infinite-dimensional optimization problem. Next, Theorem 1 states that the solution to (6) is guaranteed to be contained in an explicitly characterized finite-dimensional subspace, simplifying the optimization problem. Let  $\mathcal{H}_P$  be an RKHS of scores at  $P$  with a mean-zero PD kernel  $K_P$  on the sample space  $\mathcal{O}$ , and  $\mathbf{O} = [O_1, \dots, O_n]^\top \in \mathcal{O}^n$  be distinct data points. Define the  $n$ -dimensional subspace  $S_n := \text{span}\{k_{O_1}, \dots, k_{O_n}\} \subset \mathcal{H}$  and let  $\pi_n : \mathcal{H} \rightarrow S_n$  be the orthogonal projection onto  $S_n$ .

**Theorem 1** (RKHS representer). *Assume that  $\widetilde{\mathcal{M}}_P$  is closed under the projection onto  $S_n$ : whenever  $p(h) \in \widetilde{\mathcal{M}}_P$ , so is  $p(\pi_n h) \in \widetilde{\mathcal{M}}_P$ . Then for any loss function of the form  $L_{p,h}(O) = \tilde{L}(O, p(O), h(O))$ , the minimizer of the regularized empirical risk*

$$h^* \in \arg \min_{h \in \mathcal{H} : p(h) \in \widetilde{\mathcal{M}}_P} \mathbb{P}_n L_{p,h} + \lambda \|h\|_{\mathcal{H}}^2 \quad (7)$$

admits the following representation:  $h^* = \alpha^\top k_{\mathbf{O}} + \check{h}$ , for some  $\alpha \in \mathbb{R}^n$ , where, if  $\lambda = 0$ , then  $\check{h} \in \mathcal{H}$  is any such that  $\check{h}(O_i) = 0$  for each  $i \in [n]$ , and, if  $\lambda > 0$ , then  $\check{h} \equiv 0 \in \mathcal{H}$ .

Our choice of loss function in (6),  $L_{p,h}(O) = -\log((1 + h(O))p(O))$ , satisfies the loss function conditions of Theorem 1, and therefore admits a solution of the form  $h = \alpha^\top k_{\mathbf{O}} + \check{h}$ . Thus, Theorem 1 implies that in order to minimize (6) over the perturbations  $\widetilde{\mathcal{M}}_P$ , it suffices to consider the minimization problem over the  $n$ -dimensional submodel

$$\begin{aligned} \widetilde{\mathcal{M}}_P \supset \widetilde{\mathcal{M}}_{P,n} &:= \{p(\alpha^\top k_{\mathbf{O}}) := (1 + \alpha^\top k_{\mathbf{O}})p; \\ &\alpha \in \mathbb{R}^n \text{ such that } p(\alpha^\top k_{\mathbf{O}}) > 0\} \cap \mathcal{M}, \end{aligned} \quad (8)$$

leading to the following simpler (parametric) MLE:

$$\begin{aligned} \alpha_n(p) &:= \alpha_n(p; \lambda) := \\ &\arg \min_{\alpha \in \mathbb{R}^n : p(\alpha^\top k_{\mathbf{O}}) \in \widetilde{\mathcal{M}}_{P,n}} \mathbb{P}_n [-\log(1 + \alpha^\top k_{\mathbf{O}})p] + \lambda \alpha^\top K_{\mathbf{O}} \alpha. \end{aligned} \quad (9)$$

**Remark 1** (Regularization). *The MLE step in (6) serves as an update and stopping criterion in KDPE. It does not attempt to approximate the values of the IF. Since we are not initializing an estimation problem from scratch but rather commencing with a pre-existing estimate assumed to be of high quality, there is no necessity in MLE (6) to conform to any specific functional form or regularity, such as sparsity or other typical assumptions found in nonparametric estimation problems that might justify regularization. The*

objective of solving (6) is to update the distributional estimate in a manner that (i) maintains the consistency of the naive plugin (Assumptions 1 and 2) and (ii) satisfies the score equation (12) for every function in the RKHS (not a subset). Consequently, we have the freedom to choose any solution for (6) that fulfills requirements (i) and (ii), without imposing any additional constraints on this selection.

The proof of the asymptotic linearity and efficiency of KDPE (Theorem 2) relies on the universality of the RKHS  $\mathcal{H}_{P_n^\infty}$  rather than the span of  $\{K_{P_n^\infty}(O_1, \cdot), \dots, K_{P_n^\infty}(O_n, \cdot)\}$ .

### 3.1. Kernel Debiased Plug-in Estimator (KDPE)

Given a pre-estimate  $P_n^0$  and a PD kernel  $K$ , we

- (i) construct fluctuations  $\widetilde{\mathcal{M}}^0 \equiv \widetilde{\mathcal{M}}_{P_n^0}$  around  $P_n^0$  defined in Equation (5) by computing the zero-mean kernel  $K_{P_n^0}$  as described in Proposition 1;
- (ii) solve the regularized MLE (6) over scores  $\mathcal{H}_{P_n^0}$  via the equivalent simpler optimization (9) to find  $\alpha_n^0 \equiv \alpha_n(p_n^0)$ , which defines the MLE score  $[\alpha_n^0]^\top k_{\mathcal{O}}^0$ .

As a result, we obtain the 1<sup>st</sup>-step of our estimator  $p_n^1 = p_n^0([\alpha_n^0]^\top k_{\mathcal{O}}^0) = (1 + [\alpha_n^0]^\top k_{\mathcal{O}}^0)p_n^0$ . We define the  $\ell + 1$ <sup>th</sup>-step of KDPE recursively by

$$p_n^{\ell+1} := p_n^\ell([\alpha_n^\ell]^\top k_{\mathcal{O}}^\ell) = (1 + [\alpha_n^\ell]^\top k_{\mathcal{O}}^\ell)p_n^\ell \quad (10)$$

$$\text{and } \alpha_n^\ell \equiv \alpha_n(p_n^\ell). \quad (11)$$

This recursion terminates when  $\alpha_n^\ell = 0$ , indicating that  $p_n^\ell$  is a local optimum for the likelihood over submodel  $\widetilde{\mathcal{M}}^\ell$ . Similarly to TMLE, we assume the following about the convergence (in the iteration step  $\ell$ ) of the KDPE recursion:

**Assumption 3** (Convergence and termination at an interior point). *We assume that the  $\ell$ <sup>th</sup>-step of KDPE estimate  $p_n^\ell$  converges almost surely to some limit  $p_n^\infty$  in the interior of  $\mathcal{M}$ , and that this limiting distribution  $P_n^\infty$  is a local minimum of the regularized empirical risk (6) in the interior of the perturbation space  $\widetilde{\mathcal{M}}^\infty$ .*

Assumption 3 requires the convergence of our procedure to a limit  $P_n^\infty$ , ensuring that this limit satisfies an interior condition on 1) the full model class  $\mathcal{M}$  to ensure that our problem is regular,<sup>10</sup> and 2) the submodel  $\widetilde{\mathcal{M}}^\infty$  so that the first-order optimality conditions hold. Proposition 2 highlights a key consequence of Assumption 3:

**Proposition 2.** *Under Assumption 3, the final KDPE estimate  $P_n^\infty$  of the true distribution  $P^*$  satisfies the following estimating equations:*

$$\mathbb{P}_n h = 0 \quad \text{for all } h \in \mathcal{H}_{P_n^\infty}. \quad (12)$$

<sup>10</sup>That is, our model is locally fully nonparametric, i.e., the tangent space is  $L_0^2(P_n^\infty)$ , and our target parameter is pathwise differentiable, i.e., the conditions in Lemma 1 are satisfied.

While Proposition 2 states that KDPE solves infinitely many score equations, it guarantees that the plug-in bias,  $\mathbb{P}_n \phi_{P_n^\infty}$ , is eliminated if the influence function  $\phi_{P_n^\infty}$  falls in the RKHS  $\mathcal{H}_{P_n^\infty}$ . By choosing a universal kernel (Carmeli et al., 2010; Micchelli et al., 2006), we strengthen the Proposition 2 to achieve our desired debiasing result, under the regularity conditions below.

**Assumption 4.** *We assume that (i) RKHS  $\mathcal{H}_K$  is universal, (ii) the estimate  $P_n^\infty$  satisfies  $p^*/p_n^\infty \in L^2(P_n^\infty)$ , and (iii) the estimated influence function  $\phi_{P_n^\infty}$  is in  $L^2(P^*)$   $P^*$ -a.s.*

**Assumption 5.** *There exist an event  $\Omega$  with  $P^*(\Omega) = 1$ , and for every  $\omega \in \Omega$  there exists a sequence  $h_j(\omega) \in \mathcal{H}_{P_n^\infty(\omega)}$  indexed by  $\omega \in \Omega$  such that the following hold for all  $w \in \Omega$ : (i)  $\|h_j(\omega) - \phi_{P_n^\infty(\omega)}\|_{L^2(P^*)} \xrightarrow{j \rightarrow \infty} 0$  and (ii) there exists a  $j_0(w) \in \mathbb{N}$ ,  $j_0(w) < \infty$  such that the set of functions  $\{h_j(\omega); \omega \in \Omega, j > j_0(w)\}$  is  $P^*$ -Donsker.*

**Theorem 2** (Asymptotic linearity and efficiency of KDPE). *Let  $\psi: \mathcal{M} \rightarrow \mathbb{R}$  be a pathwise-differentiable functional of the distribution  $P$  with influence function  $\phi_P \in L_0^2(P)$  and von Mises expansion:  $\psi(\bar{P}) - \psi(P) = \int \phi_P d(P - P) + R_2(\bar{P}, P)$  for any  $\bar{P}, P \in \mathcal{M}$  (which defines the second-order remainder term  $R_2(\bar{P}, P)$ .) Under Assumptions 1 (i), 2, 3, 4, 5,  $\mathbb{P}_n \phi_{P_n^\infty} = o_{P^*}(1/\sqrt{n})$  and the KDPE estimator satisfies  $\psi(P_n^\infty) - \psi(P^*) = \mathbb{P}_n \phi_{P^*} + o_{P^*}(1/\sqrt{n})$ .*

Theorem 2 shows that KDPE attains asymptotic efficiency by converging to the optimal limiting distribution within the class of RAL estimators. To converge to this asymptotically optimal distribution, error terms that deviate from this limit distribution must be on the scale of  $o(1/\sqrt{n})$ . Thus, the definition of asymptotic efficiency (van der Vaart, 2000, Section 8) also implies that KDPE achieves local asymptotic minimax optimality with respect to the error  $\psi(P_n^\infty) - \psi(P^*)$  on the order of  $O(1/\sqrt{n})$ . The proof of Theorem 2 (Appendix B.5) requires that as  $n$  approaches infinity, the  $\sqrt{n}$ -scaled error between our true value and KDPE estimate converges to the asymptotically normal distribution implied by  $\mathbb{P}_n(\phi_{P^*})$ . The proof consists of three steps: Step 1) our uniform approximation guarantees are with respect to the true distribution  $P^*$  (as opposed to the KDPE distribution  $P_n^\infty$ ), Step 2) decompose the error of the KDPE estimate similar to that of Equation (2), and Step 3) showing that the empirical process term is of order  $o_{P^*}(1/\sqrt{n})$ .

**Remark 2.** *Assumption 5 is our novel regularity assumption on the model  $\mathcal{M}$  and functional  $\psi$  needed to control the plug-in bias term  $\mathbb{P}_n \phi_{P_n^\infty}$  of the proposed KDPE estimator. Assumptions 1-2 and conclusions derived from them are used for the one-step correction and TMLE (van der Laan and Rubin, 2006; Kennedy, 2022), and analogous conditions to Assumptions 3 and 4 are common for TMLE (van der Laan and Rubin, 2006). This is intentional: considering a similar set of assumptions directly characterizes the addi-*

tional assumption (Assumption 5) needed for KDPE and its practical advantages. Our novel assumption, Assumption 5, controls an additional, nonstandard empirical process term that arises in the analysis of the KDPE estimator, as shown in Appendix C.4. One trivial condition for Assumption 5 is that  $\phi_{P_n^\infty} \in \mathcal{H}_{P_n^\infty}$  for all sample paths. For estimation problems with nuisances (e.g., densities, conditional regression functions, etc.) over a compact, continuous input space  $\mathcal{O}$ , sufficient conditions for Assumption 5 are (i) the choice of the universal kernel (e.g., Gaussian Kernel) and (ii) bounded total variation of the approximating sequence  $\{h_j(\omega)\}_{j \in \mathbb{N}}$  for all  $\omega \in \Omega$  (see Dudley, 1992). While the universal approximation property for our choice of kernel is well-documented, our work is the first to apply this property to reduce first-order bias in nonparametric estimation.

## 4. Implementation/Simulation Results

In this section, we provide a detailed implementation of KDPE, and evaluate it using two simulated DGPs that occur in the causal inference literature, validating the results provided by Theorem 2. We note that these models are semi-parametric due to conditional independence constraints. In Appendix C, we provide the necessary modifications to Algorithm 1 for our updated distributions to be valid within the model  $\mathcal{M}$ . Additional details and results (e.g., bootstrap confidence intervals, computational complexity, update mechanism) are included in Appendix C.<sup>11</sup>

**Modification for KDPE** To use solvers, we work with a slightly modified algorithm, Algorithm 1. The main modifications from the baseline KDPE include the incorporation of a density bound parameter,  $c$ , and a convergence tolerance function. The constant  $c$  avoids violations of condition (ii) of Assumption 4 (square integrability of  $p^*/p_n^\infty$ ) and enables standard convex/concave optimization software for finding  $\alpha^{i+1}$ . The convergence tolerance function,  $l(\cdot, \cdot)$  acts as a metric tracking the net change in distribution between iterations.<sup>12</sup> For all simulations, we use the convergence tolerance function  $l(P, P') = \|p - p'\|_{L_2(\eta)}$ , where  $\eta$  is the counting measure.

### 4.1. DGPs and Target Parameters

We consider two common models in causal inference literature and demonstrate how to use KDPE to debias plug-in estimators for the desired target parameters. For all DGPs, we make the standard assumptions of positivity, conditional ignorability, and the stable unit treatment value assumption (SUTVA), which enable the identification of our target parameters with a causal interpretation.

<sup>11</sup>Code: <https://github.com/brianc0413/KDPE>.

<sup>12</sup>Because the objective function  $L(\cdot) = -\mathbb{P}_n \log(\cdot)$  is globally convex, when  $l(P_n^{i+1}, P_n^i)$  is small, we are close to a solution satisfying the necessary first-order conditions.

### Algorithm 1 KDPE implementation

- 1: **input:** data  $\{O_i\}_{i=1}^n$ , convergence tolerance function  $l(\cdot)$ , convergence tolerance  $\gamma$ , density bound  $c$ , regularization parameter  $\lambda$ , pre-estimate  $P_n^0$ , PD kernel  $K$ .
- 2: **initialize:** convergence tracker  $\beta(0) = \infty$  and  $\ell = 0$ .
- 3: **while**  $\beta(\ell) > \gamma$  **do**
- 4:   Construct the mean-zero kernel  $K^\ell := K_{P_n^\ell}$ .
- 5:   Update the distribution according to the log-likelihood loss function:
- 6:   Set  $\alpha_n^\ell$ :
 
$$\alpha_n^\ell = \underset{\substack{\alpha \in \mathbb{R}^n, \\ c \leq p_n^\ell(\alpha^\top k_{\mathcal{O}}^\ell) \leq 1-c}}{\arg \min} -\mathbb{P}_n \log([1 + \alpha^\top k_{\mathcal{O}}^\ell] p_n^\ell) + \lambda \alpha^\top K_{\mathcal{O}}^\ell \alpha$$
- 7:   set  $p_n^{\ell+1} = p_n^\ell([\alpha_n^\ell]^\top k_{\mathcal{O}}^\ell) = [1 + [\alpha_n^\ell]^\top k_{\mathcal{O}}^\ell] p_n^\ell$
- 8:   set  $\beta(\ell+1) = l(P_n^{\ell+1}, P_n^\ell)$ , and  $\ell = \ell + 1$ .
- 9: **end while**
- 10: **output:**  $\hat{P}_{\text{KDPE}} = P_n^\ell$ .

**DGP1: Observational Study** We define  $\mathcal{O} \equiv \mathcal{X} \times \mathcal{A} \times \mathcal{Y}$ , with baseline covariate  $X \in \mathcal{X} \equiv [0, 1]$ , binary treatment indicator  $A \in \mathcal{A} \equiv \{0, 1\}$ , and a binary outcome variable  $Y \in \mathcal{Y} \equiv \{0, 1\}$ . The DGP can be decomposed as follows:

$$P \equiv P_X \times P_{A|X} \times P_{Y|A,X}.$$

We place no restrictions on  $P_{Y|A,X}$  other than being a valid conditional density function (i.e., nonparametric). The true DGP in our simulation is given below:

$$X \sim \text{Unif}[0, 1], \quad A|X \sim \text{Bern}\left(0.5 + \frac{1}{3} \sin(50X/\pi)\right), \\ Y|A, X \sim \text{Bern}\left(0.4 + A(X - 0.3)^2 + \frac{1}{4} \sin(40X/\pi)\right).$$

**DGP2: Longitudinal Observational Study** A more complicated DGP represents a two-stage study. We define  $\mathcal{O} \equiv \mathcal{X} \times \mathcal{A}_0 \times \mathcal{L}_1 \times \mathcal{A}_1 \times \mathcal{Y}$ , with baseline covariate  $X \in \mathbb{R}$ , binary treatment indicators  $A_0, A_1$  (where  $A_i$  is the treatment at time  $i$ ), and binary outcome variables  $L_1, Y$ . The time ordering of the variables is given by:  $W \rightarrow A_0 \rightarrow L_1 \rightarrow A_1 \rightarrow Y$ . As before, the DGP can be decomposed as follows:

$$P \equiv P_X \times P_{A_0|X} \times P_{L_1|A_0,X} \times P_{A_1|L_1,A_0,X} \times P_{Y|A_1,L_1,A_0,X}$$

which have the following distribution:

$$X \sim \text{Unif}[0, 8], \quad A_0 \sim \text{Bern}(0.5) \\ L_1 = \mathbf{1}[3 + A_0 - 0.75X + \epsilon_1 > 0] \\ A_1 = \mathbf{1}[\text{expit}(-3 + 0.5X + 0.4 * L_1) \geq 0.3] \\ Y = \mathbf{1}[\text{expit}(X - 3.5 - 0.3A_0 - 0.5L_1 - 0.5A_1 + \epsilon_2) \geq 0.5],$$

where  $\epsilon_1, \epsilon_2$  are i.i.d. noise with distribution  $N(0, 1)$ .

**Target Parameters** The target parameters that we consider here are functions of the mean potential outcome under a specific treatment policy. For DGP1, we denote the mean potential outcome as:

$$\mu_a(P) \equiv P_X[P_{Y|A,X}[Y|A = a, X = x]].$$

For DGP2, we consider the mean potential outcome for a fixed treatment across all time points ( $A_1 = A_0 = a$ ):

$$\mu_a(P) \equiv P_X[P_{L_1|A=a,X}[P_{Y|A_1,L_1,A_0,X}[Y|A_1 = a, L_1 = l_1, A_0 = a, X = x]]].$$

The target parameters considered in this study are:

- avg. treatment effect,  $\psi_{\text{ATE}}(P) = \mu_1(P) - \mu_0(P)$ ,
- relative risk,  $\psi_{\text{RR}}(P) = \mu_1(P)/\mu_0(P)$ ,
- odds ratio,  $\psi_{\text{OR}}(P) = \frac{\mu_1(P)/(1-\mu_1(P))}{\mu_0(P)/(1-\mu_0(P))}$ .

## 4.2. Simulation Setup

**Pre-estimate Initialization** For both DGPs, we initialize the distribution of baseline covariate  $X$  as  $P_n^0(X) = \mathbb{P}_n(X)$  and use the package `SuperLearner` (van der Laan et al., 2007) to estimate the remaining conditional distributions, obtaining  $P_n^0$ . For our testing of highly adaptive lasso (HAL, van der Laan, 2017), we use the implementation of HAL provided by package `hal9001` (Coyle et al., 2022).<sup>13</sup>

**Baseline Methods for Comparison** We compute these estimators with four different plug-in methods: 1) our proposed method KDPE, 2) TMLE, 3) HAL, and 4) a baseline using the biased super-learned distribution, SL (van der Laan et al., 2007), which is our pre-estimate. TMLE has access to the functional form of the IF of each target parameter and is known for its efficient limiting distribution. TMLE serves as a benchmark for examining the asymptotic efficiency of KDPE. We expect KDPE to perform no worse than TMLE. HAL serves as a benchmark for the empirical performance of KDPE against an existing method with similar aims: HAL also solves a rich class of score / estimating equations and can be used as a debaised plug-in for many different target parameters. In contrast, comparisons against SL illustrate the effect of the KDPE approach on the distribution of our estimates, highlighting the need for using a debiasing approach. For DGP2, we use the longitudinal version of TMLE, denoted as LTMLE (van der Laan and Gruber, 2012) as our baseline. For (L)TMLE, we learn a distinct data-generating distribution for each target parameter, i.e.,  $\hat{P}_{\text{TMLE}}^{\text{RR}}, \hat{P}_{\text{TMLE}}^{\text{OR}}, \hat{P}_{\text{TMLE}}^{\text{ATE}}$ , to debias the estimate. KDPE,

<sup>13</sup>We defer additional details on hyperparameter selection and tuning to Appendix C.

	Method	$\psi_{\text{ATE}}$	$\psi_{\text{RR}}$	$\psi_{\text{OR}}$
DGP1	SL	0.0803	0.2623	0.6796
	TMLE	0.0574	0.1723	0.4059
	KDPE	0.0592	0.1752	0.4303
	HAL	0.0635	0.1987	0.4795
DGP2	SL	0.0508	0.0925	0.1555
	LTMLE	0.0731	0.1481	0.2648
	KDPE	0.0295	0.0778	0.0827
	HAL	0.0589	0.1084	0.1848

Table 1. Root Mean Squared Error (RMSE) of KDPE, (L)TMLE, HAL-MLE, and SL for DGP1, DGP2.

HAL<sub>2</sub> and SL only use a single estimated distribution  $\hat{P}_{\text{SL}}$  and  $\hat{P}_{\text{KDPE}}$ , respectively.

**Simulation Settings** In all experiments, we use 300 samples ( $n = 300$ ) and the mean-zero Gaussian kernel. For both DGP1 and DGP2, we set  $c = 0.001$ . For DGP1, we use 450 simulations, with  $\lambda = 0, \gamma = 0.002$ . For DGP2, we use 350 simulations with  $\lambda = 15, \gamma = 0.0001$ .<sup>14</sup> The choice of hyperparameters is discussed in Appendix C.

## 4.3. Comparing Empirical Performance

**DGP1 verifies the results of Theorem 2** Our simulations for DGP1 (Figures 1 and 5) demonstrate two distinct results of Theorem 2: (i) KDPE is roughly equidistant to the efficient limiting normal distribution as TMLE, and (ii) KDPE works well as a plug-in distribution for many pathwise differentiable parameters.

As an example, the limiting distribution of  $\psi_{\text{ATE}}(P^*)$  is given by  $N(0.1233, 0.0025)$  for DGP1. For DGP1, the distribution of  $\psi_{\text{ATE}}(\hat{P}_{\text{SL}})$  in Figure 1 (and  $\psi_{\text{RR}}(\hat{P}_{\text{SL}}), \psi_{\text{OR}}(\hat{P}_{\text{SL}})$ , shown in Figure 5 of Appendix C) highlight the necessity of debiasing. Without correction, the naive plug-in estimator deviates significantly from the limiting distribution, while all three debiasing approaches,  $\psi_{\text{ATE}}(\hat{P}_{\text{KDPE}}), \psi_{\text{ATE}}(\hat{P}_{\text{HAL}})$  and  $\psi_{\text{ATE}}(\hat{P}_{\text{TMLE}}^{\text{ATE}})$ , achieve similar distributions to the limiting normal distribution. Despite no knowledge of the influence function when performing debiasing, the simulated distribution of  $\psi_{\text{ATE}}(\hat{P}_{\text{KDPE}})$  is as close to the efficient limiting distribution as  $\psi_{\text{ATE}}(\hat{P}_{\text{TMLE}}^{\text{ATE}})$ , and closer to the efficient limiting distribution than  $\psi_{\text{ATE}}(\hat{P}_{\text{HAL}})$ . The root-mean-squared error results provided in Table 1 corroborate this result, where the RMSE of KDPE is closer to that of TMLE than HAL.

The results for DGP1 in Figure 5 verifies the second consequence of Theorem 2: for all pathwise differentiable parameters that satisfy the stated regularity conditions,  $\hat{P}_{\text{KDPE}}$  is debaised as a plug-in estimator. Across all three target

<sup>14</sup>Because larger values of  $\lambda$  result in smaller changes, we recommend setting smaller  $\gamma$  for larger  $\lambda$ .






Figure 1. Simulated distributions of  $\hat{\psi}_{\text{ATE}}$  compared to their asymptotic distributions. TMLE distribution in the second row corresponds to LTMLE for DGP2.

estimands, the KDPE estimate  $\psi(\hat{P}_{\text{KDPE}})$  performs similarly (both in terms of similarity to the limiting distribution and RMSE) compared with TMLE, which uses a distinct targeted plug-in ( $\hat{P}_{\text{TMLE}}^{\text{ATE}}, \hat{P}_{\text{TMLE}}^{\text{RR}}, \hat{P}_{\text{TMLE}}^{\text{OR}}$ ) for each parameter.

**DGP2 demonstrates that KDPE may provide improved finite sample performance.** DGP2 provides an example where KDPE vastly outperforms TMLE and HAL for fixed sample size  $n = 300$ , despite TMLE using the influence functions for each estimand. KDPE significantly outperforms TMLE, HAL, and the naive SL plug-in estimates both in terms of RMSE, as shown in Table 1 and convergence to the asymptotic limiting distribution, as shown in Figure 1 for  $\psi_{\text{ATE}}$ . The resulting plug-in estimates from KDPE demonstrate much smaller variance (Figures 1 and 5) than SL, HAL, or TMLE for all target parameters, and is centered correctly for  $\psi_{\text{ATE}}, \psi_{\text{OR}}$ . The performance of KDPE for  $n = 300$  indicates that for smaller sample sizes, KDPE may outperform a more targeted approach; in contrast, the IF-based approach, LTMLE, requires a much larger sample size ( $n \approx 800$ ) to reach the correct limiting distribution.

## 5. Conclusions and Practical Benefits

In this paper, we introduce KDPE, a new method for debiasing plug-in estimators that (i) does not require the IF as input, and (ii) works as a general plug-in for any target parameters that satisfy our regularity conditions.

To see the practical benefits of KDPE, let us compare the required work from the analyst (1) using existing IF-based methods, (2) using existing IF-free methods, and (3) using KDPE. An analyst may initially want to estimate  $k$  parameters from the distribution,  $\{\psi_i(P^*)\}_{i=1}^k$  from the data. Us-

ing IF-based methods (DML, TMLE, one-step), the analyst is required to derive (or find) the IF for each target parameter  $\psi_i$ , and uses the IF to debias each parameter, resulting in  $k$  analytical derivations and  $k$  runs of the debiasing algorithm. Using finite-differencing approaches, the analyst approximates the plug-in bias term  $\frac{1}{n} \sum_{i=1}^n \phi_{\hat{P}}(O_i)$ , and uses the one-step correction, resulting in  $k$  runs of the finite differencing algorithm. Using KDPE, the analyst simply obtains a distributional estimate  $\hat{P}$  from running KDPE **once**, and uses this as a plug-in estimator for the  $k$  parameters, resulting in estimates  $[\psi_i(\hat{P})]_{i=1}^k$ . After estimating the  $k$  parameters, the analyst wishes to do a follow-up analysis by estimating an additional  $w$  parameters, indexed  $[\psi_i(P^*)]_{i=k+1}^{k+w}$ . Using IF-based or finite-differencing approaches, the analyst must derive or approximate the IF  $w$  additional times, and perform  $w$  additional runs of their debiasing algorithm. Using KDPE, we simply use the same  $\hat{P}$  obtained previously as a plug-in to obtain estimates  $[\psi_i(\hat{P})]_{i=k+1}^{k+w}$ .

This example demonstrates the benefits of KDPE; it (1) avoids any computation/derivation of the IF for all  $k + w$  parameters, and (2) obtains a single debiased plug-in distribution  $\hat{P}$  that can be used to obtain debiasing for many target parameters. In contrast with existing approaches, KDPE only needs to be run once to obtain a general plug-in distribution for all  $k + w$  parameters, and greatly simplifies the estimation process for many different parameters of interest.

Future work includes (i) analyzing the effect of RKHS-norm regularization on the convergence and efficiency, (ii) finding simple heuristics for setting hyperparameters  $c, \gamma, \lambda$ , (iii) investigating how to apply sample-splitting and whether it can relax our assumptions, and (iv) verifying the validity of bootstrap.

## Acknowledgements

This work was supported by the NDSEG fellowship and Harvard Data Science Initiative Fellowship. Additionally, this research was completed while Kyra Gan was a postdoc at Harvard, and was supported by NIH grants P41EB028242 and P50DA054039. The authors would like to thank Susan Murphy for her early feedback.

## Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

## References

- Agrawal, R., Witty, S., Zane, A., and Bingham, E. (2024). Automated efficient estimation using monte carlo efficient influence functions. *arXiv preprint arXiv:2403.00158*.
- Bickel, P. J., Klaassen, C., Ritov, Y., and Wellner, J. (1998). *Efficient and adaptive estimation for semiparametric models*. Springer.
- Carmeli, C., De Vito, E., Toigo, A., and Umanitá, V. (2010). Vector valued reproducing kernel hilbert spaces and universality. *Analysis and Applications*, 8(01):19–61.
- Carone, M., Luedtke, A. R., and van der Laan, M. J. (2018). Toward computerized efficient estimation in infinite-dimensional models. *Journal of the American Statistical Association*.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., and Newey, W. (2017). Double/debiased/neyman machine learning of treatment effects. *American Economic Review*, 107(5):261–265.
- Chernozhukov, V., Newey, W. K., and Singh, R. (2022). Automatic debiased machine learning of causal and structural effects. *Econometrica*, 90(3):967–1027.
- Coyle, J. R., Hejazi, N. S., Phillips, R. V., van der Laan, L., and van der Laan, M. J. (2022). hal9001: The scalable highly adaptive lasso. R package version 0.4.2.
- Dudley, R. M. (1992). Frechet differentiability, p-variation and uniform donsker classes. *The Annals of Probability*, 20(4):1968–1982.
- Fu, A., Narasimhan, B., and Boyd, S. (2020). CVXR: An R package for disciplined convex optimization. *Journal of Statistical Software*, 94(14):1–34.
- Fukumizu, K. (2009). Exponential manifold by reproducing kernel hilbert spaces. *Algebraic and Geometric Methods in Statistics*, page 291–306.
- Gruber, S. and van der Laan, M. J. (2012). tml: An R package for targeted maximum likelihood estimation. *Journal of Statistical Software*, 51(13):1–35.
- Hines, O., Dukes, O., Diaz-Ordaz, K., and Vansteelandt, S. (2022). Demystifying statistical learning based on efficient influence functions. *The American Statistician*, 76(3):292–304.
- Jordan, M., Wang, Y., and Zhou, A. (2022). Empirical gateaux derivatives for causal inference. *Advances in Neural Information Processing Systems*, 35:8512–8525.
- Kennedy, E. H. (2022). Semiparametric doubly robust targeted double machine learning: a review. *arXiv preprint arXiv:2203.06469*.
- Koshevnik, Y. A. and Levit, B. Y. (1976). On a non-parametric analogue of the information matrix. *Teoriya Veroyatnostei i ee Primeneniya*, 21(4):759–774.
- Kotz, S. and Johnson, N. (1991). *Breakthroughs in Statistics*. Number v. 1-2 in Breakthroughs in Statistics. Springer.
- 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.
- Micchelli, C. A., Xu, Y., and Zhang, H. (2006). Universal kernels. *Journal of Machine Learning Research*, 7(12).
- Newey, W. K. (1990). Semiparametric efficiency bounds. *Journal of Applied Econometrics*, 5(2):99–135.
- Robins, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9):1393–1512.
- Smale, S. and Zhou, D.-X. (2007). Learning theory estimates via integral operators and their approximations. *Constructive Approximation*, 26(2):153–172.
- van der Laan, M. (2017). A Generally Efficient Targeted Minimum Loss Based Estimator based on the Highly Adaptive Lasso. *The International Journal of Biostatistics*, 13(2).
- van der Laan, M. and Gruber, S. (2016). One-step targeted minimum loss-based estimation based on universal least favorable one-dimensional submodels. *The International Journal of Biostatistics*, 12:351–378.
- van der Laan, M., Polley, E., and Hubbard, A. (2007). Super learner. Technical Report Working Paper 222., U.C. Berkeley Division of Biostatistics Working Paper Series.

- van der Laan, M. and Rose, S. (2018). *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*. Springer Science and Business Media.
- van der Laan, M. J., Benkeser, D., and Cai, W. (2022). Efficient estimation of pathwise differentiable target parameters with the undersmoothed highly adaptive lasso. *The International Journal of Biostatistics*.
- van der Laan, M. J. and Gruber, S. (2012). Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The International Journal of Biostatistics*, 8(1).
- van der Laan, M. J. and Rubin, D. (2006). Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1).
- van der Vaart, A. (1991). On differentiable functionals. *The Annals of Statistics*, pages 178–204.
- van der Vaart, A. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- Zheng, W. and van der Laan, M. (2010). Asymptotic theory for cross-validated targeted maximum likelihood estimation. Technical Report Working Paper 273., U.C. Berkeley Division of Biostatistics Working Paper Series.

## Appendix

### A. Notation

$[n]$	set of integers $1, \dots, n$
$\mathcal{O}$	Sample space, with $\mathcal{O} \subset \mathbb{R}^d$ and Borel $\sigma$ -algebra $\mathcal{B}$ ;
$\mathcal{M}$	Statistical model, collection of probability measures on $(\mathcal{O}, \mathcal{B})$ ;
$\psi(P)$	Target parameter functional, mapping from $\mathcal{M} \rightarrow \mathbb{R}$
$\phi_P$	Influence function, mapping from $\mathcal{O} \rightarrow \mathbb{R}$
$Pf, P[f]$	Expectation functional $Pf = P[f] = \int f dP$
$\mathbb{P}_n, \mathbb{P}_n f$	Empirical distribution function
$L_0^2(P)$	The set of measurable functions $\{f : \mathcal{O} \rightarrow \mathbb{R} : Pf = 0, Pf^2 < \infty\}$
$P^*$	(Unknown) data generating distribution;
$o_{P^*}(1/\sqrt{n})$	Represents a sequence of random variables that converges to zero in probability with respect to probability measure $P^*$ , at a rate that is faster than $1/\sqrt{n}$ .
$K, \mathcal{H}_K$	The PSD kernel and associated RKHS with kernel function $K(\cdot, \cdot)$
$k_x$	The function $k_x := K(x, \cdot) : \mathcal{X} \rightarrow \mathbb{R}$ , for $x$ in support $\mathcal{X}$
$k_{\mathbf{x}}$	Vectorized version of function $k_x$ , where $k_{\mathbf{x}} := [k_{x_1}, \dots, k_{x_n}]^T$
$K_{\mathbf{x}\mathbf{y}}, K_{\mathbf{x}}$	Matrix of values $K_{\mathbf{x}} := [K(x_i, x_j)]_{i,j} \in \mathbb{R}^{n \times n}$ , $K_{\mathbf{x}\mathbf{y}} := [K(x_i, y_j)]_{i,j} \in \mathbb{R}^{n \times m}$
$K_P, \mathcal{H}_P$	The mean-zero kernel (with respect to $P$ ) and associated RKHS
$\mathcal{C}_0(\mathcal{X})$	The space of continuous functions that vanish at infinity with the uniform norm
$\mathcal{C}_0$	A kernel $K$ is $\mathcal{C}_0$ provided that $\mathcal{H}_K$ is a subspace of $\mathcal{C}_0(\mathcal{X})$
$\widetilde{\mathcal{M}}_P$	The collection of perturbations $\widetilde{\mathcal{M}}_P := \{(1+h)p \mid h \in \mathcal{H}_P \text{ such that } (1+h)p > 0\} \cap \mathcal{M}$
$\widetilde{\mathcal{M}}_{P,n}$	The set of $n$ -dimensional parametric submodels, defined as follows:
	$\widetilde{\mathcal{M}}_{P,n} := \{p_{\alpha} := p(\alpha) := (1 + \alpha^\top k_{\mathbf{O}})p ; \alpha \in \mathbb{R}^n \text{ such that } p(\alpha) > 0\} \cap \mathcal{M}$
$\widetilde{\mathcal{M}}^\ell$	The collection of perturbations $\widetilde{\mathcal{M}}_{P_n^l}$ corresponding to $\ell$ -th iteration's distribution, $P_n^l$
$p(\epsilon h), p_{\epsilon,h}$	1-dimensional parametric model $(1 + \epsilon h)p$
$P_{\epsilon,h}$	The distribution corresponding to the density implied by the 1-dimensional parametric model $(1 + \epsilon h)p$
$p(h), p_h$	RKHS-based model $(1 + h)p$ , where $h \in \mathcal{H}_P$
$L_{p,h}$	Log-likelihood loss, given by $L_{p,h} \equiv -\log(p_h)$
$D_P \psi$	Derivative functional (assumed to be continuous, linear) that maps $L_0^2(P) \rightarrow \mathbb{R}$ , $\frac{d}{d\epsilon} \psi(P_{\epsilon,h}) _{\epsilon=0} = D_P \psi[h]$ for all $h \in L_0^2(P)$
$\pi_S$	The orthogonal projection operator $\mathcal{H} \rightarrow S$ , projects a function $h \in \mathcal{H}$ onto a set of functions $S$
$\alpha_n^i$	The optimal perturbation for objective function $L(\widehat{p}_n(\alpha, p), \lambda)$ at iteration $i$ for KDPE
$P_n^i, p_n^i$	Estimated data distribution of TMLE/KDPE at iteration $i$
$P_n^\infty, p_n^\infty$	Limiting/final estimate for TMLE/KDPE
$\widehat{P}$	Generic estimated distribution, subscripts/superscript added accordingly
$\widehat{\psi}_{\text{param.}}$	Estimated target parameter value $\psi_{\text{param.}}(\widehat{P})$ using plug-in distribution $\widehat{P}$ , subscripts/superscript added accordingly
$Q(P)$	Relevant portion of distribution (e.g. $\{Q_a(X)\}_{a \in \{0,1\}} = \{P[Y = 1 X, A = a]\}_{a \in \{0,1\}}$ ).
$\Omega$	An event (or set of events) such that has probability 1 under $P^*$ , i.e. $P^*(\Omega) = 1$ . The standard example is the set of all possible $n$ -sample realizations $\{\{O_i(w)\}_{i=1}^n : w \in \Omega\}$ from distribution $P^*$ .
$f(w), P(w)$	The random function / distribution, indexed by events $w \in \Omega$ . Often, we use $w \in \Omega$ to index an $n$ -sample realization, $\{O_i(w)\}_{i=1}^n$ , from DGP $P^*$ .

## B. KDPE Theory

### B.1. Additional Details Regarding $P_{\epsilon,h}$ and $\widetilde{\mathcal{M}}_P$ .

**Validity of Linear Submodels.** The linear submodel  $P_{\epsilon,h}$  is indeed a valid submodel (i) for all  $h \in L_0^2(P)$  and (ii) for  $|\epsilon| \leq \epsilon_{h,P}$ , where  $\epsilon_{h,P}$  is an upper bound on the magnitude of  $\epsilon$  depending on  $h$  and  $P$ . For example,  $h = p'/p - 1$ , the path connecting  $P$  and another unique (up to the dominating measure) distribution  $P' \in \mathcal{M}$ , has  $0 \leq \epsilon \leq \epsilon_{h,P} = 1$ , which corresponds to  $p_{\epsilon,h} = (1 - \epsilon)p + \epsilon p'$ ; for  $h' = h/2$ , our bound  $\epsilon_{h',P}$  becomes 2. Because we are in the nonparametric setting, the only constraint on our model class  $\mathcal{M}$  is that we remain a valid density, and thus our linear paths  $P_{\epsilon,h}$  are fully contained in the model class  $\mathcal{M}$ .

**Additional Comments Regarding  $\widetilde{\mathcal{M}}_P$ .** We note that  $\widetilde{\mathcal{M}}_P$  will never be empty in the nonparametric case. Intuitively, Eq. 5 simply restricts the feasible scores  $h \in L_0^2(P)$  to a subset of scores  $h \in \mathcal{H}_P$ , which is dense in  $L_0^2(P)$  and therefore nonempty for any  $P$ . At each  $P$ , the nonparametric tangent space is  $L_0^2(P)$ , which means that all  $h \in L_0^2(P)$  are feasible directions. In other words, one can move at least an infinitesimal amount in each  $h \in L_0^2(P)$  and remain within the model class. In our bounded kernel setting with universal (e.g. Gaussian) kernels,  $\mathcal{H}_P$  is a dense subset of  $L_0^2(P)$ . This implies for each  $P$  and  $h \in \mathcal{H}_P$ ,  $p(h)$  is a valid submodel in  $\mathcal{M}$ , given a correct magnitude (i.e. obeys  $\epsilon_{h,P}$  bounds implied by  $\mathcal{M}$ , where  $h$  is rescaled to be unit norm). Thus, the restriction imposed in Eq. 5 corresponds to bounding the maximum magnitude  $\epsilon$  we can move for each unit-norm  $h \in \mathcal{H}_P$  (i.e.  $\epsilon_{h,P}$ ), such that we remain a valid density.

### B.2. Proof of Proposition 1

*Proof of Proposition 1.* Because  $K$  is bounded, the expectation functional  $h \mapsto Ph$  is continuous on  $\mathcal{H}_K$ . Next, we note that  $\varphi \in \mathcal{H}_K$  and it is the Riesz representer of the expectation functional (we have shown this elsewhere and remark here that the rigorous argument considers a Riemann approximation to the expectation integral and weakly convergent subsequences). It follows that  $h \in \mathcal{H}_P$  if and only if  $\langle h, \varphi \rangle_K = 0$ , and from this we can deduce the projection operator onto  $\mathcal{H}_P$  and the reproducing kernel of that subspace.

Finally, given  $h \in L_0^2(P)$ , let  $\{g_n\} \subset \mathcal{H}_K$  be an approximating sequence of  $h$  in  $\mathcal{H}_K$ . Let  $g_n = \tilde{g}_n + \mathbb{E}_P[g_n]\varphi/\mathbb{E}_{P \times P}[K]$  be the RKHS orthogonal decomposition of  $g_n$  into  $\mathcal{H}_P$  and  $\mathcal{H}_P^\perp$ , and note that this operation is continuous on  $L^2(P)$ . By continuity of expectation in  $L^2(P)$ ,  $\mathbb{E}_P[g_n] \rightarrow 0$  in  $L^2(P)$ . It follows that the orthogonal part of  $g_n$  vanishes in  $L^2(P)$  and  $\tilde{g}_n \rightarrow h$  in  $L^2(P)$ . Thus,  $\mathcal{H}_P$  is dense in  $L_0^2(P)$ .  $\square$

### B.3. Proof of Theorem 1

Let  $L(Q(P)) : \mathcal{X} \rightarrow \mathbb{R}$  be a general loss function that identifies the relevant part,  $Q(P)$ , for the target parameter of interest  $\psi(P) = \psi(Q(P))$ , satisfying the property that  $\mathbb{E}_P L(Q(P)) \leq \mathbb{E}_P L(Q')$  for every possible  $Q'$ . The log-likelihood  $L(P) = -\log p$  is a special case of such a loss function. The risk associated with the log-likelihood identifies the entire distribution  $P$  based on the properties of the Kullback–Leibler divergence.

*Proof of Theorem 1.* For  $h \in \mathcal{H}$  consider the orthogonal decomposition  $h = h_S + h_S^\perp$  into the closed subspaces  $S_n$  and  $S_n^\perp$ . From the reproducing property it follows that  $h(O_i) = \langle h_S + h_S^\perp, K_{O_i} \rangle = h_S(O_i)$  and  $h_S^\perp(O_i) = 0$  for every  $i \in [n]$ . This means that inputs  $h$  and  $h_S$  yield the same value of the empirical risk  $\sum_i \tilde{L}(O_i)$  in the objective function of (6). Furthermore, we have assumed that if  $h$  is admissible in the optimization problem (6), then so is  $h_S$ . The claim follows by noting that  $\|h\|_{\mathcal{H}}^2 = \|h_S\|_{\mathcal{H}}^2 + \|h_S^\perp\|_{\mathcal{H}}^2$ .  $\square$

### B.4. Proof of Proposition 2

**Remark 3.** Assumption 3 assumes that our algorithm converges and terminates in finite time. It is important to highlight that this assumption is commonly adopted in the TMLE literature (van der Laan and Rubin, 2006). As demonstrated in Table 2, our algorithm, Algorithm 1, converges effectively in practice, typically requiring only a few iterations. While the regularization of the perturbation size (i.e., the norm of  $h$ ) could potentially ensure finite convergence, it is beyond the scope of this paper and remains a topic for future research.

*Proof of Proposition 2.* This follows from the first order condition to (6) which must hold at the interior local minimum. First, consider the variations at  $p_n^\infty$  in  $\widetilde{\mathcal{M}}_{P_n^\infty}$  along the  $n$  directions (scores) given by  $K_{P_n^\infty}(O_1, \cdot), \dots, K_{P_n^\infty}(O_n, \cdot)$  and

observe:

$$\begin{aligned}
 0 &= \frac{\partial}{\partial \alpha_j} \Big|_{\alpha=0} \sum_i -\log \left( 1 + \sum_j \alpha_j (K_{P_n^\infty})(O_j, O_i) \right) p_n^\infty(O_i) + \lambda_K \alpha^\top K_{\mathbf{O}} \alpha \\
 &= \left[ - \sum_i \frac{K_{P_n^\infty}(O_j, O_i) p_n^\infty(O_i)}{(1 + \sum_j \alpha_j (K_{P_n^\infty})(O_j, O_i)) p_n^\infty(O_i)} + \lambda_K (K_{\mathbf{O}} \alpha)_j \right] \Big|_{\alpha=0} \\
 &= - \sum_i K_{P_n^\infty}(O_j, O_i).
 \end{aligned}$$

This shows that the estimating equation holds for each of these scores, and therefore their linear combinations. Second, by Theorem 1, any  $h \in \mathcal{H}_{P_n^\infty}$  that is orthogonal to these scores must have  $h(O_i) = 0$  for all  $i \in [n]$  and therefore satisfy Equation (12). From these two observations, we conclude the proposition.  $\square$

### B.5. Proof of Theorem 2

**Remark 4.** Assumptions 2-3 and conclusions derived from them are standard in the TMLE literature (van der Laan and Rubin, 2006; Kennedy, 2022). Assumption 4 is mostly technical and ensures that the set of scores that satisfy the estimating equation (12) is sufficiently rich to achieve debiasing for the target parameter by approximating its efficient influence curve estimating equation asymptotically.

Assumption 5 is our main regularity assumption on the model  $\mathcal{M}$  and functional  $\psi$  needed to control the plug-in bias term  $\mathbb{P}_n \phi_{\hat{p}}$  of the proposed KDPE estimator. While this assumption provides a clear intuition for the debiasing mechanism of KDPE, it is not strictly required for the method to work. Less extreme conditions on the regularity of  $\mathcal{F}$ , such as  $C^\alpha$  or Sobolev norm bounds (see van der Vaart 2000, Ch 19 and RKHS approximation theory Smale and Zhou 2007), can imply the Donsker property in S5.

The Donsker assumption S2 has been successfully relaxed for related estimators in the literature (Zheng and van der Laan, 2010; Kennedy, 2022; Chernozhukov et al., 2017) via sample-splitting. Our initial analysis of DKPE crucially relies on the Donsker assumption in the argument for asymptotic debiasing (based on S5). We leave the analysis of KDPE with sample-splitting (as done in Zheng and van der Laan 2010) to future work.

Recall that the proof consists of three steps: Step 1) our uniform approximation guarantees are with respect to the true distribution  $P^*$  (as opposed to the KDPE distribution  $P_n^\infty$ ), and Step 2) decompose the error of the KDPE estimate similar to that of Equation (2), and Step 3) showing that the empirical process term is of order  $o_{P^*}(1/\sqrt{n})$ .

**Step 1** is achieved by establishing the existence of a sequence in  $\mathcal{H}_{P_n^\infty}$  that approximates  $\phi_{P_n^\infty}$  arbitrarily well with respect to the  $L^2(P^*)$  norm via change of measure (via Assumption 4, square integrability of density ratios, and the density of  $\mathcal{H}_{P_n^\infty}$  in  $L_0^2(P_n^\infty)$ .) In **Step 2**, by using the result from Step 1, the gradient condition of Assumption 3, and the expansion in Lemma 1, we obtain that the error of the KDPE estimate can be decomposed as four terms: (i) the empirical average of  $\phi_{P^*}$ , which gives us our desired result, (ii) a standard empirical process term across other estimation methods (e.g. TMLE, DML, one-step), (iii) a standard second-order remainder, and (iv) an empirical process term that involves the approximating sequence  $h_j$ . Terms (ii) and (iii) are controlled by Assumptions 1 and 2 respectively such that they are  $o_{P^*}(1/\sqrt{n})$ . In **Step 3**, we use our novel assumption, Assumption 5, to establish that the nonstandard term (iv) such that it is of order  $o_{P^*}(1/\sqrt{n})$ . Thus, terms (ii-iv) vanish at the appropriate rate, resulting in the conclusion of Theorem 2.

*Proof of Theorem 2.* Because  $\phi_{P_n^\infty} \in L_0^2(P_n^\infty)$ , and the RKHS  $\mathcal{H}_{P_n^\infty}$  is dense in  $L_0^2(P_n^\infty)$ , there must exist a sequence  $\{h_j\}_{j \in \mathbb{N}} \subset \mathcal{H}_{P_n^\infty}$ , where for all  $\epsilon > 0$ , there exists a  $j(\epsilon) \in \mathbb{N}$  such that for all  $j > j(\epsilon)$ ,  $\|h_j - \phi_{P_n^\infty}\|_{L^2(P_n^\infty)} < \epsilon$ .

To get the desired convergence rates with respect to true measure  $P^*$ , we first note that under Assumption 4,  $|P_n^\infty(p^*/p_n^\infty)|$  is bounded by some constant  $\eta(\omega)$  for sample paths  $\omega \in \Omega$  by Holder's inequality:

$$P_n^\infty(p^*/p_n^\infty) \leq P_n^\infty[1]^{1/2} \times \underbrace{\|p^*/p_n^\infty\|_{L^2(P_n^\infty)}}_{< \infty} = \eta(\omega)$$

Assumption 4 also enables change of measure by a dominating measure  $\lambda$ . By the Cauchy-Schwartz inequality,

$$\begin{aligned} \|h_j - \phi_{P_n^\infty}\|_{L^2(P^*)}^2 &= \int (h_j - \phi_{P_n^\infty})^2 dP^* = \int (h_j - \phi_{P_n^\infty})^2 \left(\frac{p_n^\infty}{P_n^\infty}\right) p^* d\lambda(O) \\ &\leq \int \frac{p_n^*}{P_n^\infty} dP_n^\infty \times \int (h_j - \phi_{P_n^\infty})^2 dP_n^\infty(O) \\ &\leq \eta(\omega) \|h_j - \phi_{P_n^\infty}\|_{L^2(P_n^\infty)}^2 \end{aligned}$$

By setting  $j > j(\epsilon/\sqrt{\eta(\omega)})$ , this leads to an equivalent result under  $L^2(P^*)$  norm (rather than  $L^2(P_n^\infty)$  norm):

$$\|h_j - \phi_{P_n^\infty}\|_{L^2(P^*)} \leq \sqrt{\eta(\omega)} \|h_j - \phi_{P_n^\infty}\|_{L^2(P_n^\infty)} < \epsilon.$$

Thus, for all  $\epsilon > 0$ , there exists  $j(\epsilon/\sqrt{\eta(\omega)}) \in \mathbb{N}$  such that  $j > j(\epsilon/\sqrt{\eta(\omega)})$  implies  $\|h_j - \phi_{P_n^\infty}\|_{L^2(P^*)} < \epsilon$ . An immediate consequence is another useful result for our analysis. By a simple application of Holder's inequality,

$$\begin{aligned} |P^*(h_j - \phi_{P_n^\infty})| &\leq P^*[1 \times |h_j - \phi_{P_n^\infty}|] \\ &\leq P^*(|h_j - \phi_{P_n^\infty}|^2)^{1/2} \times P^*(1^2)^{1/2} \quad (\text{Holder's Inequality for } p = q = 2.) \\ &= \|h_j - \phi_{P_n^\infty}\|_{L^2(P^*)} < \epsilon. \end{aligned}$$

These two results show that for any sample path  $\omega \in \Omega$ , there exists a sequence  $\{h_j\}_{P_n^\infty} \subset \mathcal{H}_{P_n^\infty}$ , such that for all  $\epsilon > 0$ , there exists a  $j(\epsilon/\sqrt{\eta(\omega)}) \in \mathbb{N}$  where  $j > j(\epsilon/\sqrt{\eta(\omega)})$  implies the following:

$$\text{(A)} : \|h_j - \phi_{P_n^\infty}\|_{L^2(P^*)} < \epsilon, \quad \text{(B)} : |P^*(h_j - \phi_{P_n^\infty})| < \epsilon.$$

Under the conditions of Assumption 3 and 4, we analyze our estimator using the expansion in Lemma 1:

$$\psi(P_n^\infty) - \psi(P^*) = \mathbb{P}_n \phi_{P^*} - \underbrace{\mathbb{P}_n \phi_{P_n^\infty} + (\mathbb{P}_n - P^*)[\phi_{P_n^\infty} - \phi_{P^*}]}_{(a) = \text{"double empirical process term"}} + \underbrace{R_2(P_n^\infty, P^*)}_{(b) = \text{second-order remainder}}. \quad (13)$$

Under Assumption 2, term (b) is  $o_{P^*}(1/\sqrt{n})$ , and is therefore asymptotically negligible. It remains to show that term (a) vanishes appropriately. Let  $\{j^*(\omega, n)\}_{n \in \mathbb{N}}$  be the index sequence such that  $\epsilon = o_{P^*}(1/\sqrt{n})$  for sample path  $\omega$ . Under Assumption 3,  $\mathbb{P}_n h = 0$  for every  $h \in \mathcal{H}_{P_n^\infty}$ , giving us the following equality:

$$(a) = -\mathbb{P}_n(\phi_{P_n^\infty}) + (\mathbb{P}_n - P^*)(\phi_{P_n^\infty} - \phi_{P^*}) = -\mathbb{P}_n(-h_{j^*} + \phi_{P_n^\infty}) + (\mathbb{P}_n - P^*)(\phi_{P_n^\infty} - \phi_{P^*}).$$

Given property (B), we can rewrite the "double empirical process term" (a) as follows:

$$\begin{aligned} (a) &= -\mathbb{P}_n(\phi_{P_n^\infty} - h_{j^*}) + (\mathbb{P}_n - P^*)(\phi_{P_n^\infty} - \phi_{P^*}) \\ &= -\mathbb{P}_n(\phi_{P_n^\infty} - h_{j^*}) + (\mathbb{P}_n - P^*)(\phi_{P_n^\infty} - \phi_{P^*}) + P^*(\phi_{P_n^\infty} - h_{j^*}) - \underbrace{P^*(\phi_{P_n^\infty} - h_{j^*})}_{=o_{P^*}(1/\sqrt{n})} \\ &= -[\mathbb{P}_n(\phi_{P_n^\infty}) - P^*(\phi_{P_n^\infty})] + [\mathbb{P}_n(h_{j^*}) - P^*(h_{j^*})] + (\mathbb{P}_n - P^*)(\phi_{P_n^\infty} - \phi_{P^*}) + o_{P^*}(1/\sqrt{n}) \\ &= (\mathbb{P}_n - P^*)(h_{j^*} - \phi_{P^*}) + o_{P^*}(1/\sqrt{n}) \\ &= \underbrace{(\mathbb{P}_n - P^*)(h_{j^*} - \phi_{P_n^\infty})}_{(c)} + \underbrace{(\mathbb{P}_n - P^*)(\phi_{P_n^\infty} - \phi_{P^*})}_{(d)} + o_{P^*}(1/\sqrt{n}). \end{aligned}$$

Term (d) =  $o_{P^*}(1/\sqrt{n})$  by Assumption 1, and term (c) =  $o_{P^*}(1/\sqrt{n})$  by Assumption 5 and Property (A) via Lemma 19.24 of van der Vaart, 2000. Thus, term (a) =  $o_{P^*}(1/\sqrt{n})$ . We conclude that  $\psi_{P_n^\infty} - \psi(P^*) = \mathbb{P}_n \phi_{P^*} + o_{P^*}(1/\sqrt{n})$ .  $\square$

## C. Simulation Details and Additional Empirical Results

This section provides pseudocode for the bootstrap procedure, empirical justification for our choice of hyperparameters  $c, \gamma$ , and additional experiments for the bootstrap variance estimation procedure. All simulations were performed on a Dell Desktop with 11th Gen Intel Core i7, 16 GB of RAM. All code for the simulations provided in both this section and Section 4 can be found on <https://github.com/anonymous/KDPE>.

### C.1. Baseline Methods

**Initializing  $P_n^0$  for KDPE** For all simulations, we use the SuperLearner (van der Laan et al., 2007) and `tmle`, `ltml` (Gruber and van der Laan, 2012; van der Laan and Gruber, 2012) packages in R. For SuperLearner, we set our initial base learners as `SL.randomForest`, `SL.glm`, and `SL.mean`. For `tmle`, `ltml`, we use the default settings, except for setting the parameter `g_bound` to 0. For each run of the simulations, we initialize  $P_n^0$  for both TMLE and KDPE to the same initial density estimate to capture the differences in the two debiasing approaches.

**Optimization Step Details** For the optimization solver, we use the package CVXR (Fu et al., 2020) with the default settings to obtain  $\alpha^{i+1}$ , which corresponds to the parametric update step in Algorithm 1. Because the log-likelihood loss function  $L(\alpha, P) = \mathbb{P}_n \log((1 + \sum_{i=1}^n \alpha_i K_P(o_i, \cdot))p)$  is concave and all constraints are linear, our update step can be solved by existing convex optimization software. To avoid solvers, a potential future direction is to explore alternative loss functions that have known closed-form solutions, such as the mean squared error (MSE), rather than using the log-likelihood loss.

**Construction of  $K_P(o_i, \cdot)$**  To construct our  $n$ -dimensional parametric submodels, we use the mean-zero RBF kernel  $K_P(\cdot, \cdot) : \mathcal{O} \times \mathcal{O} \rightarrow \mathbb{R}$ , which takes the following form:

$$K_P(O, O') = \exp(-\|O - O'\|_2^2) - \frac{\int \exp(-\|O - s\|_2^2) dP(s) \int \exp(-\|O' - s\|_2^2) dP(s)}{\int \int \exp(-\|s - t\|_2^2) dP(s) dP(t)},$$

where  $O, O' \in [0, 1] \times \{0, 1\}^2$ . Following classical examples (van der Laan and Rose, 2018) and existing software (Gruber and van der Laan, 2012) for TMLE, we fix  $X \sim \mathbb{P}_n(x)$  and  $A \sim P_n^0(A|X)$  (the initial estimate for the propensity score function). This enables us to directly calculate the integral, rather than approximate  $K_P$ :

$$K_P(O, O') = \exp(-\|O - O'\|_2^2) - \frac{f_P(O)f_P(O')}{c_P}, \quad (14)$$

where  $f_P$  and  $c_P$  are defined as follows:

$$f_P(O') = f_P((x', a', y')) = \frac{1}{n} \sum_{i=1}^n \sum_{(a,y) \in \{0,1\}^2} P(A = a, Y = y | X = x_i) \exp(-\|(x', a', y') - (x_i, a, y)\|^2),$$

$$c_P = \frac{1}{n} \sum_{i=1}^n \sum_{a \in \{0,1\}} \sum_{y \in \{0,1\}} P(A = a, Y = y | X = x_i) f(x_i, a, y).$$

Note that this calculation is inherent to our setting of binary treatments  $A$  and outcomes  $Y$  with the empirical marginal distribution for  $X$  – other set-ups may not permit a direct calculation of the integral, and may require integral approximations. We defer those experiments to future work.

**Settings for HAL-MLE** We use the HAL implementation provided by Coyle et al., 2022. For DGP1, we estimate the conditional densities  $P(Y|A, X)$ ; for DGP2, we estimate conditional densities  $P(Y|A_1, L_1, A_0, X)$  and  $P(L_1|A_0, X)$ . All estimation with HAL is done with base settings, using the link function setting `family = "binomial"` to respect the  $[0, 1]$  bounds we make on our data-generating process.

### C.2. Restricted Tangent Space Modifications

The following update steps for DGP1 and DGP2 provide the projection step required for the restricted tangent spaces tested in our simulations. These updates provide explicit examples of conditional tangent space updates, where one is only allowed to update a specific component of the initial DGP estimate, but imposes non-parametric assumptions on this component. A simple example of this restricted model setting is treatment effect estimation with known propensity scores,



but no restrictions on the conditional outcome function. Another example is offline estimation of a proposed policy's value, where the treatment assignment policy is specified by the user, but no restrictions are made on the the reward distribution. In the exposition of KDPE, we assume that we update the entire density  $P_n^\ell(O)$ , and therefore  $L_0^2(P_n^\ell)$  is the correct tangent space under our nonparametric assumptions. However, in our simulations, we only aim to perturb the conditional density functions (i.e.  $P(Y|A, X)$  in DGP1,  $P(L_1|A_0, X)$ ,  $P(Y|A_1, L_1, A_0, X)$  in DGP2). For example, in DGP1, the scores (i.e. perturbation directions) of the conditional density function  $P(Y|A, X)$  live in the conditional/projected tangent space  $L_0^2(P_n^\ell(Y|A, X))$ , as opposed to the more general tangent space of  $L_0^2(P_n^\ell)$ .

To get the specific updates for DGP1, DGP2, we project the chosen score (i.e.,  $h \equiv [\alpha_n^\ell]^\top k_{\mathbf{O}}^\ell$ ) into the relevant tangent space. The only change in Algorithm 1 is the following step, which occurs at the end of each loop:

$$p_n^{\ell+1} = p_n^\ell([\alpha_n^\ell]^\top k_{\mathbf{O}}^\ell) = [1 + [\alpha_n^\ell]^\top k_{\mathbf{O}}^\ell] p_n^\ell.$$

For all set-ups in our simulation, the tangent space factorizes nicely, which means we simply need to project our chosen score  $h$  at each iteration. We demonstrate the necessary modifications to obtain our results.

**DGP1** Since the conditional density function  $P_n^\ell(A|X)$  does not show up in the expression of our final plug-in estimate of the target parameters (ATE, RR, OR), it suffices to consider only updating the conditional density function  $P_n^\ell(Y|A, X)$ , with distribution  $P_n^\ell(X) = \mathbb{P}_n(X)$  fixed. The tangent space  $L_0^2(P_n^\ell)$  decomposes as follows:

$$L_0^2(P_n^\ell) = L_0^2(P_n^\ell(X)) \oplus L_0^2(P_n^\ell(A|X)) \oplus L_0^2(P_n^\ell(Y|A, X)),$$

where  $L_0^2(P_n^\ell(Y|A, X))$  is space of scores (i.e., directions of change) corresponding to  $P_n^\ell(Y|A, X)$ . To obtain the component of our chosen score  $h = [\alpha_n^\ell]^\top k_{\mathbf{O}}^\ell$  relevant for  $P_n^\ell(Y|A, X)$ , we simply project this into  $L_0^2(P_n^\ell(Y|A, X))$ :

$$h' = h - P_n^\ell(h|A, X), \quad p_n^{\ell+1}(Y|A, X) = (1 + h')p_n^\ell(Y|A, X).$$

**DGP2** Similarly, since  $P(A_1|L_1, A_0, X)$  and  $P(A_0|X)$  do not show up in the expression of our final plug-in estimate of the target parameters, we do not need to update these components. We fix the marginal distribution of baseline covariates as  $P_n^0(X) = \mathbb{P}_n(X)$ . The two components relevant for our target parameters are  $P(L_1|A_0, X)$ , and  $P(Y|A_1, L_1, A_0, X)$ . The tangent space for DGP2 factorizes as follows:

$$L_0^2(P_n^\ell) = L_0^2(P_n^\ell(X)) \oplus L_0^2(P_n^\ell(A_0|X)) \oplus L_0^2(P_n^\ell(L_1|A_0, X)) \oplus L_0^2(P_n^\ell(A_1|L_1, A_0, X)) \oplus L_0^2(P_n^\ell(Y|A_1, L_1, A_0, X)),$$

where  $L_0^2(P_n^\ell(L_1|A_0, X))$  and  $L_0^2(P_n^\ell(Y|A_1, L_1, A_0, X))$  are the corresponding tangent spaces for  $P_n^\ell(L_1|A_0, X)$  and  $P_n^\ell(Y|A_1, L_1, A_0, X)$  respectively. At the end of each iteration, we project the chosen score  $h$  into the tangent space of these relevant components, and update accordingly:

$$\begin{aligned} h' &= P_n^\ell(h|L_1, A_0, X) - P_n^\ell(h|A_0, X), & p_n^{\ell+1}(L_1|A_0, X) &= (1 + h')p_n^\ell(L_1|A_0, X), \\ h'' &= h - P_n^\ell(h|A_1, L_1, A_0, X), & p_n^{\ell+1}(Y|A_1, L_1, A_0, X) &= (1 + h'')p_n^\ell(Y|A_1, L_1, A_0, X). \end{aligned}$$

The steps outlined above project the loss-minimizing / likelihood-maximizing updates from scores in  $L_0^2(P_n^\ell)$  to scores in the correct tangent spaces implied by the component in our density estimate we wish to update (i.e.  $L_0^2(P(Y|A, X))$  in DGP1,  $L_0^2(P(L_1|A_0, X))$ ,  $L_0^2(P(Y|A_1, L_1, A_0, X))$  in DGP2).

**Other Semi-Parametric Extensions.** Beyond projecting the updates (i.e. loss-minimizing perturbations) directly in our algorithm, analogous results (i.e. efficiency, asymptotic normality) for semi-parametric models can be directly achieved by our KDPE framework with the appropriate choice of kernel function. For closed, linear tangent spaces  $T_P$ , let  $H_P \equiv H \cap T_P$  be the intersection of the tangent space  $T_P$  and a universal RKHS  $H$ . The intersected space  $H_P$  itself is an RKHS, with corresponding kernel  $K_P$ . By using this kernel for KDPE, one obtains the same results for the semiparametric models. For nonparametric models,  $K_P$  was simply the kernel corresponding to the projection of  $H$  onto mean-zero functions, given in Proposition 1. In semi-parametric settings, this requires a projection of  $K$  onto  $H \cap T_P$  in the norm of the RKHS  $H$ . For general semi-parametric settings, this projection for the tangent space may have to be derived analytically. While it might seem that we have traded analytic derivation of the EIF for an analytic derivation for the appropriate RKHS kernel

function, the plug-in approach of KDPE works for *all target parameters* that satisfy our stated assumptions 1-5. In other words, KDPE replaces the analytic derivation of the IF for each parameter of interest defined on the same semi-parametric model (requires derivation for each parameter of interest, even within the same set of assumptions) with a single analytic characterization of the RKHS projection onto the tangent space of the semi-parametric model (only requires derivation for new modeling assumptions). We chose the non-parametric setting for our paper as the simplest setting to explain KDPE, and plan to explore these extensions in future work.

### C.3. Hyperparameter Testing for KDPE

**Density Bound and Convergence Criteria** In our experiments, we fixed the sample size at  $n = 300$ . We conducted tests with different values for the density bound  $c$ , with  $c \in \{0.05, 0.01, 0.002\}$ , and for the stopping criteria, we tested  $\gamma \in \{0.025, 0.005, 0.001\}$  for DGP1, with  $\lambda$  fixed at 0. For each combination of settings, we performed 100 simulations to obtain the results presented in Table 2 and Figure 2.

$c$	$\gamma$	ATE RMSE	RR RMSE	OR RMSE	Avg. Iterations
0.050	0.025	0.019	0.077	0.205	1.062
0.010	0.025	0.016	0.063	0.180	1.135
0.002	0.025	0.026	0.101	0.260	1.120
0.050	0.005	0.015	0.057	0.150	2.293
0.010	0.005	0.004	0.026	0.091	2.445
0.002	0.005	0.006	0.033	0.090	2.505
0.050	0.001	0.001	0.023	0.068	4.756
<b>0.010</b>	<b>0.001</b>	<b>0.009</b>	<b>0.040</b>	<b>0.128</b>	<b>5.329</b>
<b>0.002</b>	<b>0.001</b>	<b>0.009</b>	<b>0.033</b>	<b>0.105</b>	<b>5.802</b>

Table 2. Average RMSE and iterations until convergence for DGP 1. No regularization, i.e.,  $\lambda = 0$ .

The empirical results shown in Figure 2 coincide with the theoretical results shown in Appendix B, which are derived under the setting where we have no explicit density bound ( $c = 0$ ) and we terminate at a fixed point ( $\gamma = 0$ ). As we see in Figure 2, the simulated distribution of  $\psi_{\text{ATE}}(\hat{P}_{\text{KDPE}})$  with the smallest tested values of  $c, \gamma$  obtains a distribution most similar to limiting normal distribution. When the density bound is set to large values (e.g.,  $c = 0.05$ ), the resulting distributions become skewed and deviate significantly from the normal distribution. This occurs because termination at the boundary (i.e., is tight with respect to the density bound) prevents the distributions from satisfying the gradient conditions shown in Appendix B. When using large values for the convergence tolerance (e.g.,  $\gamma = 0.025$ ), the algorithm terminates prematurely at a distribution that deviates significantly from the fixed point  $P_n^\infty$ . Thus, to maintain the guarantees shown in Appendix B for Algorithm 1, we want the values of  $(\gamma, c)$  to be as small as possible. For DGP1 in Sections 4.3, we opt for  $c = 0.001$ ,  $\gamma = 0.001$ , values at least as small as the tested hyperparameter values. The results of this simulation naturally pose questions of how to set  $c, \gamma$  as functions of  $n$ , such that the residual bias caused by these hyperparameters is  $o_{P^*}(1/\sqrt{n})$ . This remains an open question for future work.

**Regularization Parameter  $\lambda$**  While the theoretical guarantees of KDPE hold for *any*  $\lambda \geq 0$ , the choice of  $\lambda$  has practical significance, especially in relation to the convergence tolerance parameter  $\gamma$ . Large values of  $\gamma$  (i.e., looser convergence tolerances) with heavy regularization (i.e., large  $\lambda$ ) results in premature termination. We demonstrate these results with DGP1, testing  $\lambda \in \{0, 0.001, 0.01, 0.1, 1, 10, 100\}$ <sup>15</sup>. For each  $\lambda$ , we ran 100 simulations with  $n = 300$  to determine the best choice of  $\lambda$ . As shown in Figure 3, for a fixed value of convergence tolerance  $\gamma = 0.001$ , we see that larger values of  $\lambda$  tend to result in skewed distributions and/or heavy tails. The average iterations for convergence corroborate that this is due to premature termination; the resulting step with heavy regularization results in a relatively small change in distribution and trivially satisfies our convergence criteria. For this reason, for DGP2 (which uses  $\lambda = 20$ ), we set the convergence tolerance to  $\gamma = 0.0001$ , as compared to  $\gamma = 0.001$  in the non-regularized case in DGP1. We keep the density bound for DGP2 the same ( $c = 0.001$ ). Future extensions of this work include how to set regularization parameter  $\lambda$  as a function of convergence

<sup>15</sup>Note that for  $\lambda = 0$ , which we use throughout our paper, the solver CVXR fails to obtain a solution roughly 15% of the time. For those interested in testing the code, we recommend  $\lambda = 0.001$ , which results in similar results, but avoids numerical instability issues.




Figure 2. Density histogram of  $\psi_{\text{ATE}}(\hat{P}_{\text{KDPE}})$  simulations as a function of  $c, \gamma$  for DGP1. X-axis: ATE values. Y-axis: density. Blue dotted line: the true ATE under  $P^*$ , and the y-axis indicates the count.

tolerance  $\gamma$ .

#### C.4. Additional (L)TMLE Results for DGP2

We provide additional results to show that the poor performance of TMLE in DGP2 is due to the relatively small sample size  $n = 300$ . For  $n = 800$ , Figure 4 see that the distribution of  $\hat{\psi}_{\text{TMLE}}^{\text{ATE}}$  is far closer to the limiting distribution (shown in purple) than when  $n = 300$ , as shown in the main body of our paper. This indicates that the performance of TMLE for DGP2 at  $n = 300$  is indeed due to poor finite sample performance.

#### C.5. Additional Empirical Results for Relative Risk and Odds Ratio

To show that our asymptotic normality holds for multiple target parameters  $\psi$ , we include the simulated distributions of the KDPE estimate against TMLE and SL in Figure 5, using the same hyperparameter settings described above.

#### C.6. Runtime Complexity for KDPE

Our algorithm is a meta-algorithm, and thus the computational complexity is dependent on different factors (integration method, optimization software, number of iterations, etc.). We provide computational complexity results per iteration in the case of Algorithm 1. The set-up necessary to formulate the optimization problem requires us to obtain  $K_P(O_i, O_j)$  (Eq. 20) for all  $i, j \in [n]$ , and occurs the cost  $\mathcal{O}(n^2)$ :

- Computing  $\exp(-\|O_i - O_j\|^2)$  for all  $i, j \in [n]$ :  $\mathcal{O}(n^2)$




Figure 3. (a): Average number of iterations, with maximum and minimum number of iterations for each  $\lambda$  tested ;(b): Distribution of  $\hat{\psi}_{KDPE}^{ATE}$  in DGP1 for different  $\lambda$ ,  $c = 0.002$ ,  $\gamma = 0.0001$ . Blue line denotes true value. 100 simulations for each setting of  $\lambda$ .




Figure 4. Distribution of  $\hat{\psi}_{LTMLE}^{ATE}$  for  $n = 800$ , using LTMLE Package for DGP2. Purple shows correct limiting distribution.




Figure 5. Simulated distributions of  $\hat{\psi}_{RR}, \hat{\psi}_{OR}$ . First row corresponds to DGP1, and second row to DGP2. Red line denotes true value of the target parameter.

- Computing  $f_P(O_i)$  for all  $i \in [n]$ :  $\mathcal{O}(4n^2)$
- Computing  $c_P$ :  $\mathcal{O}(4n)$

The optimization problem for KDPE involves  $n$  regressors, and therefore the computational complexity of this method is  $\mathcal{O}(n^2 + k(n))$  per iteration, where  $k(n)$  is the complexity for the solver.

In contrast, TMLE (as introduced in the paper) occurs an  $\mathcal{O}(n)$  cost when calculating the influence function in each iteration, and involves only 1 regressor, resulting in a complexity of  $\mathcal{O}(n + k(1))$  in each iteration.

The difference in cost per iteration (quadratic, as opposed to linear in  $n$ ), can be seen as the price to pay for a plug-in distribution that works with all parameters satisfying our regularity assumptions. Our empirical results corroborate these scaling results. The median time-per-iteration in DGP1 for 150 samples and 300 samples is 0.22 and 0.91 respectively, indicating a roughly  $\mathcal{O}(n^2)$  scaling for our method. For DGP1, we average between 5-6 iterations before termination, using the hyperparameters in Section 4. For DGP2, we average 2 iterations before termination. Across our simulations for DGP1, the average runtime for our method is 5.8 seconds, with 5.8 iterations on average. TMLE, as implemented in (van der Laan and Gruber, 2016), uses the one-step method, which performs the MLE optimization in one iteration using logistic regression. Because this version of TMLE involves no iterations (same time complexity as logistic regression), it is a poor comparison with KDPE in terms of time complexity. The results of TMLE are included as a standard of comparison for the distribution of our estimator, rather than its time complexity.

### C.7. Bootstrap Algorithm

We use the classical bootstrap procedure (Kotz and Johnson, 1991), i.e., sampling  $n$  observations from our observed data  $\{o_i\}_{i=1}^n$  with replacement, to estimate the variance of the estimator. We formally state this procedure in Algorithm 2.

### C.8. Bootstrap Intervals for Inference

Table 3 reports average estimated variance, average length of 95% intervals, and the coverage across 237 simulations of  $n = 300$  for DGP1. For comparison, we include the results of inference for TMLE, which estimates the variance  $\widehat{\text{Var}}(\hat{\psi}) = \mathbb{P}_n[(\hat{\phi}_{\hat{P}_{\text{TMLE}}}^{\text{target}})^2]$  using knowledge of the (efficient) IF. The baselines for both TMLE and KDPE (in parentheses)

**Algorithm 2** Bootstrap Confidence Interval for KDPE

- 1: **Input:**  $\{O_i\}_{i=1}^n$ , convergence tolerance function  $l(\cdot)$ , convergence tolerance  $\gamma$ , density bound  $c$ , and bootstrap iterations  $m$ .
- 2: Obtain distribution  $\widehat{P}_{\text{KDPE}} = \text{KDPE}(\{O_i\}_{i=1}^n, l(\cdot), \gamma, c)$ .
- 3: Set the estimate  $\widehat{\psi} = \psi(\widehat{P}_{\text{KDPE}})$ .
- 4: **for**  $j \in \{1, \dots, m\}$  **do**
- 5: Sample with replacement  $\{O_i^*\}_{i=1}^n$  from sample  $\{O_i\}_{i=1}^n$ , and remove duplicates.
- 6: Set the estimate  $\widehat{\psi}_j = \psi(\text{KDPE}(\{O_i^*\}_{i=1}^n, l(\cdot), \gamma, c))$
- 7: **end for**
- 8: Calculate the mean estimate  $\bar{\psi} = \sum_{j=1}^m \widehat{\psi}_j / m$ .
- 9: Estimate the variance  $\widehat{S} = \sum_{j=1}^m (\widehat{\psi}_j - \bar{\psi})^2 / (m - 1)$ .
- 10: Return  $\left[ \widehat{\psi} \pm \Phi^{-1}(1 - \alpha) \sqrt{\widehat{S}} \right]$ .

	Parameter	$\widehat{\text{Var}}(\widehat{\psi}_{\text{target}})$	Avg. Length (95% CI)	Coverage (95% CI)
KDPE	$\psi_{\text{ATE}}$	0.00427 (0.00287)	0.252	0.945
	$\psi_{\text{RR}}$	0.04562 (0.02555)	0.826	0.975
	$\psi_{\text{OR}}$	0.30303 (0.15322)	2.107	0.970
TMLE	$\psi_{\text{ATE}}$	0.00192 (0.00296)	0.172	0.903
	$\psi_{\text{RR}}$	0.01483 (0.02837)	0.510	0.903
	$\psi_{\text{OR}}$	0.07016 (0.14981)	1.230	0.907

Table 3. Estimated Variance and Confidence Intervals of  $\widehat{\psi}_{\text{ATE}}, \widehat{\psi}_{\text{RR}}, \widehat{\psi}_{\text{OR}}$  of KDPE and TMLE for DGP1 across 237 simulations, with 100 bootstrap iterations each for KDPE.

are the sample variances of the simulated distributions in Section 4.3, and provided next to average estimated variance. The KDPE-bootstrap variance estimates overestimate the variance for all target parameters relative to the KDPE baseline; the inflated variance estimates lead to conservative confidence intervals for the bootstrapped KDPE estimates with over-coverage, as shown in the coverage results of Table 3. By contrast, TMLE’s IF-based variance estimate underestimates the variance for all target parameters relative to the TMLE baseline, leading to undercoverage. These preliminary results suggest that bootstrap interval calibration is future direction for improvement, and are intended as a starting point for future investigation. While these results imply that our bootstrap-estimated confidence intervals are conservative, more extensive testing is required to empirically validate these results. A large theoretical gap also remains for confirming that our bootstrap-estimated variances are consistent asymptotically, and providing an analysis for Algorithm 2 under KDPE is another natural step towards our goal of fully computerized inference. Lastly, as shown in Table 3 and Algorithm 2, we only use the bootstrap samples to estimate the variance of our estimator and use a normal approximation to construct confidence intervals. Alternative methods for constructing confidence intervals include directly taking the quantiles of the bootstrap estimates and/or subsampling while bootstrapping, which we do not explore in this project.