

One-Step Estimation of Differentiable Hilbert-Valued Parameters

Alex Luedtke and Incheoul Chung

Department of Statistics, University of Washington

November 24, 2025

Abstract

We present estimators for smooth Hilbert-valued parameters, where smoothness is characterized by a pathwise differentiability condition. When the parameter space is a reproducing kernel Hilbert space, we provide a means to obtain efficient, root- n rate estimators and corresponding confidence sets. These estimators correspond to generalizations of cross-fitted one-step estimators based on Hilbert-valued efficient influence functions. We give theoretical guarantees even when arbitrary estimators of nuisance functions are used, including those based on machine learning techniques. We show that these results naturally extend to Hilbert spaces that lack a reproducing kernel, as long as the parameter has an efficient influence function. However, we also uncover the unfortunate fact that, when there is no reproducing kernel, many interesting parameters fail to have an efficient influence function, even though they are pathwise differentiable. To handle these cases, we propose a regularized one-step estimator and associated confidence sets. We also show that pathwise differentiability, which is a central requirement of our approach, holds in many cases. Specifically, we provide multiple examples of pathwise differentiable parameters and develop corresponding estimators and confidence sets. Among these examples, four are particularly relevant to ongoing research by the causal inference community: the counterfactual density function, dose-response function, conditional average treatment effect function, and counterfactual kernel mean embedding.

1 Introduction

There has been much recent work on combining tools from semiparametric efficiency and machine learning to estimate finite-dimensional parameters (Baiardi and Naghi, 2021; Kennedy, 2022; Hines et al., 2022). These works often focus on pathwise differentiable parameters, which are characterized by their smoothness along regular univariate submodels of the statistical model (Pfanzagl, 1990; van der Vaart, 1991; Bickel et al., 1993). When a finite-dimensional parameter is pathwise differentiable, it also has an efficient influence function (EIF), which corresponds to the Riesz representation of its pathwise derivative. Efficient influence functions are the critical ingredient used to define various estimation strategies, such as those based on one-step estimation (Pfanzagl, 1982; Newey and McFadden, 1994), estimating equations (van der Laan et al., 2003; Tsiatis, 2006), targeted learning (van der Laan and Rubin, 2006; van der Laan et al., 2011), and double machine learning (Chernozhukov et al., 2017, 2018). When paired with cross-fitting (Schick, 1986; Klaassen, 1987), these frameworks yield asymptotically efficient estimators provided the nuisance functions are estimated well enough as the sample size n grows to make a certain remainder term negligible. Often, this amounts to requiring an $n^{-1/4}$ -rate condition, which will most plausibly hold if the nuisance functions are estimated flexibly.

Another line of research has focused on leveraging machine learning tools to estimate function-valued parameters, such as the causal dose-response function (Díaz and van der Laan, 2013), counterfactual

density function (Kennedy et al., 2021), and conditional average treatment effect function (Nie and Wager, 2021). Possibly owing to the wealth of available methods for estimating real-valued functionals, many of these works have focused on the evaluation of these functions at a point. As has been noted in van der Laan et al. (2018) and Chernozhukov et al. (2018), the resulting point evaluations tend not to be pathwise differentiable except in trivial cases (e.g., when the data are discrete). To overcome this challenge, kernel-smoothed approximations of the function evaluation parameter have been considered in those two works and others (Colangelo and Lee, 2020; Luedtke and Wu, 2020; Chernozhukov et al., 2021; Jung et al., 2021), and local polynomial approximations have also been introduced for several parameters (Kennedy et al., 2017; Takatsu and Westling, 2022; Kennedy et al., 2022). These smoothed approximations tend to yield pathwise differentiable parameters, which enables the use of one-step estimators. Slower-than- $n^{1/2}$ convergence rates are typically attained because the fineness of the approximation must improve with sample size. The guarantees provided for these estimators tend to be pointwise in nature. Given that pointwise convergence does not generally imply norm convergence or uniform convergence without additional regularity conditions, these pointwise-based estimators usually only facilitate inference for the evaluation of the unknown function at one or finitely many points, rather than for the entire function.

Some works have focused on estimating unknown function-valued parameters in a norm sense. Many of these works incorporate objects from semiparametric efficiency theory. For example, in the context of conditional average treatment effect estimation, risk functions have been developed (van der Laan, 2006; Luedtke and van der Laan, 2016; Nie and Wager, 2021), and Kennedy (2020) develops rate-of-convergence guarantees for the corresponding empirical risk minimizers. These estimators incorporate (weighted) variants of the EIF of the marginal average treatment effect in their construction. As further examples, EIFs have been used to construct norm-convergent estimators of the counterfactual density function (Kennedy et al., 2021) and dose-response function (Takatsu and Westling, 2022). A drawback to these approaches to estimating function-valued parameters is that, to date, it has seemed that a new estimator must be derived and new regularity conditions established for each new parameter considered. Others have presented general approaches to learning unknown functions based on empirical risk minimization, where the population risk depends on unknown nuisance functions that can be orthogonalized by conducting statistical learning using an efficient estimator of the risk function as an objective function (van der Laan and Dudoit, 2003; Foster and Syrgkanis, 2019). When the population regret takes the form of a squared norm, these approaches provide a means to derive estimators with norm-convergence guarantees. However, unlike standard approaches such as one-step-estimation that are used for estimating finite-dimensional quantities, these methods do not appear to easily lend themselves to the construction of confidence sets for the unknown functions. Instead, the available approaches to construct confidence sets rely on approaches that are generally distinct from those used for estimating the function, such as building them using higher-order influence functions (Robins et al., 2008), a restricted score test (Hudson et al., 2021), or a maximum mean discrepancy (MMD) criterion (Luedtke et al., 2019).

In this work, we establish that the one-step estimation methodology can be extended to estimate and make inference about pathwise differentiable parameters that take values in a Hilbert space. This pathwise differentiability condition turns out to be quite reasonable for many function-valued parameters of current interest. Indeed, we show that all of the parameters mentioned earlier in this Introduction satisfy it under regularity conditions. This is true in spite of the fact that these Hilbert-valued parameters are not pathwise differentiable when composed with an evaluation map.

The notion of pathwise differentiability that we focus on in this work is that studied in some early literature on semiparametric efficiency theory, which defined pathwise differentiability and EIFs not just for finite-dimensional parameters, but for general Banach-valued parameters (van der Vaart and Wellner, 1989; van der Vaart, 1991; page 179 of Bickel et al., 1993). Since all Hilbert spaces are Banach spaces, their definitions apply in our case, as do some useful results that they present, such as a convolution theorem.

Nevertheless, existing works did not provide any examples of how to evaluate the pathwise differentiability of infinite-dimensional Hilbert-valued parameters — for example, see [van der Vaart \(1991\)](#) and Chapter 5.3 of [Bickel et al. \(1993\)](#), whose infinite-dimensional examples all pertain to parameters taking values in a Banach space equipped with the uniform norm. Since they are not even pathwise differentiable at a point, none of the aforementioned function-valued parameters are pathwise differentiable in such a Banach space. Previous works also do not indicate whether or how the pathwise differentiability of an infinite-dimensional Hilbert-valued parameter can be used to facilitate estimation or inference, whether via the one-step estimation methodology or otherwise. While a brief, one-paragraph sketch was given on page 405 of [Bickel et al. \(1993\)](#) suggesting that providing a general, efficient $n^{-1/2}$ -rate estimation framework may be difficult for infinite-dimensional spaces, this sketch only discusses a single example where $n^{-1/2}$ -rate estimation may not even be possible. Moreover, neither that work, nor any subsequent ones, appear to evaluate whether leveraging the pathwise differentiability of a Hilbert-valued parameter would be useful for constructing a performant, but slower than $n^{-1/2}$ -rate, estimator, or for constructing a confidence set.

The main contributions of this work are as follows:

1. We characterize the EIF of a pathwise differentiable Hilbert-valued parameter, when it exists, and provide a means to obtain a regularized version thereof, when it does not.
2. We construct one-step estimators using (possibly regularized) EIFs. Any method can be used to estimate the needed nuisance functions provided it converges at a suitable rate.
3. We provide root- n -rate weak convergence and efficiency guarantees for these estimators, when an EIF exists, and slower rate-of-convergence guarantees, when one does not.
4. We show how to construct asymptotically-valid confidence sets for the Hilbert-valued estimands. These confidence sets take different forms depending on whether an EIF exists.
5. We study our framework in examples of current interest to the causal inference community and establish the pathwise differentiability of several more traditional parameters.

When the estimand is a function, our confidence sets will contain it with a specified probability. Thus, if the aim is to infer about the whole function, our methodology is likely preferable to pointwise approaches. To accomplish the last point above, we derive a general lemma that facilitates the evaluation of the pathwise differentiability of Hilbert-valued parameters. Finally, we have conducted a simulation study to evaluate the proposed approach. All proofs can be found in the appendix.

2 Pathwise differentiability in Hilbert spaces and constructing estimators

2.1 Notation

We work on a Polish space $(\mathcal{Z}, \mathbf{B})$ with a collection of distributions \mathcal{P} , which we refer to as the model. Let $Z_1, Z_2, \dots, Z_n \sim P_0$ be an independent and identically distributed (iid) sample from a distribution $P_0 \in \mathcal{P}$, and let P_n denote the corresponding empirical distribution. Let $\hat{P}_n \in \mathcal{P}$ be an estimate of P_0 . To ease notation, for now we consider a sample splitting approach wherein \hat{P}_n is fitted using an iid sample that is independent of Z_1, Z_2, \dots, Z_n ; in Section 2.5, we describe the case where cross-fitting is used ([Schick, 1986](#); [Klaassen, 1987](#)), which is our preferred approach. For Q a signed measure on $(\mathcal{Z}, \mathbf{B})$ and a measurable function $f : \mathcal{Z} \rightarrow \mathbb{R}$, we use the shorthand $Qf := \int f dQ$. For any object indexed by P_0 ,

we will abbreviate the notation by replacing ‘ P_0 ’ by ‘0’; for example, we will write f_0 rather than f_{P_0} . Similarly, we will replace ‘ \widehat{P}_n ’ by ‘ n ’ and write f_n rather than $f_{\widehat{P}_n}$.

All Hilbert spaces mentioned in this paper are real Hilbert spaces. For a measure μ on a measurable space (\mathcal{X}, Σ) , we write $L^2(\mu)$ to denote the Hilbert space of μ -a.s. equivalence classes of $\mathcal{X} \rightarrow \mathbb{R}$ functions equipped with inner product $\langle f, g \rangle_{L^2(\mu)} := \int f g d\mu$. If $\mathcal{X} \subset \mathbb{R}^d$, μ is the Lebesgue measure, and Σ is the Borel σ -algebra on \mathbb{R}^d , we will sometimes write $L^2(\mathcal{X})$ instead of $L^2(\mu)$. In what follows \mathcal{V} denotes a generic Hilbert space. We let $\|\cdot\|_{\mathcal{V}}$ and $\langle \cdot, \cdot \rangle_{\mathcal{V}}$ denote the norm and inner product associated with \mathcal{V} . The space $L^2(P; \mathcal{V})$ is the Hilbert space containing all Bochner measurable functions $f : \mathcal{Z} \rightarrow \mathcal{V}$ such that

$$\|f\|_{L^2(P; \mathcal{V})} := \left(\int \|f(z)\|_{\mathcal{V}}^2 P(dz) \right)^{1/2} < \infty.$$

The operator norm of a linear functional $f : \mathcal{V} \rightarrow \mathbb{R}$ is defined as $\|f\|_{\text{op}} := \inf\{c \geq 0 : |f(v)| \leq c\|v\|_{\mathcal{V}} \text{ for all } v \in \mathcal{V}\}$. If \mathcal{W} is a closed subspace of \mathcal{V} , then let $\Pi_{\mathcal{V}}[h \mid \mathcal{W}]$ denote the orthogonal projection of h to \mathcal{W} . We let ℓ^2 denote the space of all square-summable sequences and $\|b\|_{\ell^2} := [\sum_{k=1}^{\infty} b_k^2]^{1/2}$. We also let $[0, 1]^{\mathbb{N}}$ denote the space of all $[0, 1]$ -valued sequences. To avoid having to use different notation to treat finite- and infinite-dimensional Hilbert spaces, throughout we use the convention that, if \mathcal{V} is finite-dimensional, then we call $(v_k)_{k=1}^{\infty}$ an orthonormal basis of \mathcal{V} if $(v_k)_{k=1}^{\dim(\mathcal{V})}$ is an orthonormal system that spans \mathcal{V} and $v_k = 0$ for all $k > \dim(\mathcal{V})$.

2.2 Pathwise differentiability in Hilbert spaces

We start with a brief review of important definitions that can be used to characterize the smoothness of a Hilbert-valued parameter. These definitions are adapted from those given in (Bickel et al., 1993) for more general Banach-valued parameter settings. The subsequent parts of this section will involve developing estimators for our more specialized, but understudied, setting, where we heavily leverage the availability of an inner product in our Hilbert parameter space.

Let \mathcal{P} be a collection of distributions defined on a common Polish space $(\mathcal{Z}, \mathbf{B})$, which we refer to as the model. Suppose that the model is dominated by a σ -finite measure λ . A submodel $\{P_{\epsilon} : \epsilon \in [0, \delta)\} \subset \mathcal{P}$ is said to be quadratic mean differentiable at P if and only if there exists a score function $s \in L^2(P)$ such that

$$\left\| p_{\epsilon}^{1/2} - p^{1/2} - \epsilon s p^{1/2} / 2 \right\|_{L^2(\lambda)} = o(\epsilon), \quad (1)$$

where, for $\epsilon \geq 0$, $p_{\epsilon}^{1/2} = \sqrt{\frac{dP_{\epsilon}}{d\lambda}}$ and $p^{1/2} = \sqrt{\frac{dP}{d\lambda}}$. Let $\mathcal{P}(P, \mathcal{P}, s)$ refer to the set of quadratic mean differentiable submodels at P with score function s . The set $\{s \in L^2(P) : \mathcal{P}(P, \mathcal{P}, s) \neq \emptyset\}$ is called the tangent set, and its closed linear span is called the tangent space of \mathcal{P} at P , denoted by $\dot{\mathcal{P}}_P$. For all $s \in \dot{\mathcal{P}}_P$, $Ps = \int s dP = 0$. We let $L_0^2(P) := \{h \in L^2(P) : Ph = 0\}$, which is the largest possible tangent space at P . Any model with this tangent space at all distributions P it contains is referred to as locally nonparametric.

Let \mathcal{H} be a set known as the action space and $\nu : \mathcal{P} \rightarrow \mathcal{H}$ a parameter whose value is to be estimated. Throughout we assume that \mathcal{H} is a real separable Hilbert space. The parameter ν is said to be pathwise differentiable at P if and only if there exists a continuous linear operator $\dot{\nu}_P : \dot{\mathcal{P}}_P \rightarrow \mathcal{H}$ such that, for all $\{P_{\epsilon} : \epsilon \in [0, \delta)\} \in \mathcal{P}(P, \mathcal{P}, s)$,

$$\|\nu(P_{\epsilon}) - \nu(P) - \epsilon \dot{\nu}_P(s)\|_{\mathcal{H}} = o(\epsilon). \quad (2)$$

The operator $\dot{\nu}_P$ is called the local parameter of ν at P and its Hermitian adjoint, denoted by $\dot{\nu}_P^* : \mathcal{H} \rightarrow \dot{\mathcal{P}}_P$,

is referred to as the efficient influence operator. The image of the local parameter $\dot{\nu}_P$, denoted by $\dot{\mathcal{H}}_P$, is a closed subspace of \mathcal{H} that is referred to as the local parameter space. Throughout we equip $\dot{\mathcal{H}}_P$ with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$, so that $\dot{\mathcal{H}}_P$ is itself a Hilbert space. The efficient influence operator can be shown to only depend on its argument through its projection onto the local parameter space, in the sense that $\dot{\nu}_P^*(h) = \dot{\nu}_P^*(\Pi_{\mathcal{H}}[h \mid \dot{\mathcal{H}}_P])$ for all $h \in \mathcal{H}$. At times in this work, we will consider pointwise evaluations of the efficient influence operator of the form $\dot{\nu}_P^*(h)(z)$. When doing so, we always assume that suitably ‘nice’ elements of the P -a.s. equivalence classes defined by the elements $\dot{\nu}_P^*(h)$ of $L^2(P)$ are used to define these evaluations. In particular, we select these elements so that the efficient influence process, which we define as $\{\dot{\nu}_P^*(h) : h \in \mathcal{H}\}$, is a separable stochastic process, in the sense that there exists a countable dense subset \mathcal{H}' of \mathcal{H} and a P -probability one subset \mathcal{Z}' of \mathcal{Z} such that, for all $h \in \mathcal{H}$ and $z \in \mathcal{Z}'$, there exists an \mathcal{H}' -valued sequence $(h_j)_{j=1}^{\infty}$ that converges to h and satisfies $\dot{\nu}_P^*(h_j)(z) \rightarrow \dot{\nu}_P^*(h)(z)$ as $j \rightarrow \infty$.

Analogous to the case for Euclidean parameters, in some semiparametric models it may be natural to describe a Hilbert-valued parameter as the restriction of a parameter defined on a larger, possibly nonparametric, model. If the true parameter lies in a model $\mathcal{P}' \subset \mathcal{P}$ with tangent space $\dot{\mathcal{P}}'_P$ and $\nu : \mathcal{P} \rightarrow \mathcal{H}$ is a parameter defined on \mathcal{P} , then the restriction $\nu|_{\mathcal{P}'}$ has local parameter $\dot{\nu}|_{\dot{\mathcal{P}}'_P}$ and efficient influence operator $h \mapsto \Pi_{L^2(P)}[\dot{\nu}_P^*(h) \mid \dot{\mathcal{P}}'_P]$. Armed with this fact, results can easily be transferred from a larger nonparametric model to a semiparametric model provided the form of the projection operator $\Pi_{L^2(P)}[\cdot \mid \dot{\mathcal{P}}'_P]$ is known. As a simple example, we may have that $\nu(P)(\cdot) = E_P[Y \mid X = \cdot]$ for each P in a nonparametric model \mathcal{P} , and the model $\mathcal{P}' = \{P \in \mathcal{P} : \text{var}_P(Y) = 1\}$ may reflect knowledge that the variance of an outcome Y is 1. The form of the local parameter and efficient influence operator of ν relative to \mathcal{P} are given in Example 5 in the appendix when \mathcal{H} is an L^2 space, and the form of the projection onto $\dot{\mathcal{P}}'_P$ is given in Example 3.2.3 of [Bickel et al. \(1993\)](#).

The parameter ν is said to have an EIF $\phi_P : \mathcal{Z} \rightarrow \mathcal{H}$ when there exists a P -probability-one set \mathcal{Z}' such that

$$\dot{\nu}_P^*(h)(z) = \langle h, \phi_P(z) \rangle_{\mathcal{H}} \quad \text{for all } (h, z) \in \mathcal{H} \times \mathcal{Z}'. \quad (3)$$

By the Riesz representation theorem, ν has an EIF if and only if $\dot{\nu}_P^*(\cdot)(z) : \mathcal{H} \rightarrow \mathbb{R}$ is a bounded linear functional P -almost surely; in those cases, $\phi_P(z)$ is P -a.s. equal to the Riesz representation of $\dot{\nu}_P^*(\cdot)(z)$. The fact that $\dot{\nu}_P^*(h) = \dot{\nu}_P^*(\Pi_{\mathcal{H}}[h \mid \dot{\mathcal{H}}_P])$ implies that the image of ϕ_P is necessarily contained in $\dot{\mathcal{H}}_P$. When $\mathcal{H} = \mathbb{R}^d$, the EIF of ν at P takes the form $\phi_P(z) = (\dot{\nu}_P^*(e_t)(z))_{t=1}^d$, where $\{e_t\}_{t=1}^d$ is the standard basis. To our knowledge, the existence and form of this object have not previously been studied in infinite-dimensional Hilbert spaces. Given that knowing the form of the EIF readily facilitates the construction of estimators in finite-dimensional settings, this appears to constitute an important gap in the literature. We therefore focus the remainder of this section on studying the existence of EIFs in Hilbert spaces and providing ways to construct estimators based on EIFs, when we can show they exist, or imitations thereof, when we cannot.

The cases where EIFs do not exist become particularly salient in Section 3, where we demonstrate through examples that, for several interesting L^2 -valued parameters, $\dot{\nu}_P^*(\cdot)(z) : \mathcal{H} \rightarrow \mathbb{R}$ depends on a point evaluation functional, and so is not bounded P -almost surely. Nevertheless, even when ν does not have an EIF, we will show in Section 5.2 that it is always possible to define an injective transformation of ν that has one. Consequently, the procedure we shall present to construct confidence sets for parameters with EIFs can also be used to construct them for those without one: first construct a confidence set for the transformation of $\nu(P_0)$, and then invert it to obtain one for $\nu(P_0)$.

Before proceeding, we note that inefficient influence operators can be defined when the model is semiparametric at P , in that $\dot{\mathcal{P}}_P$ is a strict subspace of $L_0^2(P)$. Under a condition akin to (3), inefficient influence functions can also be defined. To streamline presentation, we defer the presentation of these

objects and their use for constructing estimators to Appendix H. There, we also show that inefficient influence functions can only exist if an EIF exists.

2.3 One-step estimation based on the efficient influence function

On the one hand, if \mathcal{H} is finite-dimensional, then the EIF can be used to construct what is known as a one-step estimator, which is known to be efficient under conditions (Pfanzagl, 1982). This estimator takes the form $\nu(\widehat{P}_n) + P_n\phi_n$, where $\widehat{P}_n \in \mathcal{P}$ is an initial estimate of the data-generating distribution P_0 and we recall the convention that $\phi_n := \phi_{\widehat{P}_n}$. On the other hand, if \mathcal{H} is infinite-dimensional, then previously studied one-step estimators cannot be applied. In this section, we provide a natural means to extend the one-step estimation framework to infinite-dimensional settings. Similarly to the one-step estimator in finite-dimensional settings, this one-step estimator takes the form $\widehat{\nu}_n := \nu(\widehat{P}_n) + P_n\phi_n$ for an \mathcal{H} -valued EIF ϕ_n . This estimator is applicable whenever ν has an EIF at \widehat{P}_n with P_0 -probability one. We will see that, under conditions that include that ν also has an EIF ϕ_0 at P_0 , $\widehat{\nu}_n$ is an asymptotically linear estimator of $\nu(P_0)$ with influence function ϕ_0 , in the sense that

$$\widehat{\nu}_n - \nu(P_0) = \frac{1}{n} \sum_{i=1}^n \phi_0(Z_i) + o_p(n^{-1/2}), \quad (4)$$

where throughout we let Hilbert-valued quantities of the form $o_p(n^{-a})$ denote terms whose Hilbert norm goes to zero in probability even after being multiplied by n^a . We will be especially interested in cases where $n^{1/2}[\widehat{\nu}_n - \nu(P_0)]$ will converge weakly to a tight random element, since this can be used to facilitate the construction of confidence sets for $\nu(P_0)$ (see Section 4.2). To be able to apply a central limit theorem to establish the weak convergence of $n^{1/2}[\widehat{\nu}_n - \nu(P_0)]$, it suffices that ϕ_0 is P_0 -Bochner square integrable, in the sense that $\|\phi_0\|_{L^2(P_0; \mathcal{H})} < \infty$ (Example 1.8.5 of van der Vaart and Wellner, 1996). Therefore, we will focus on settings where $\phi_0 \in L^2(P_0; \mathcal{H})$.

We begin by establishing the existence and form of the EIF at a generic $P \in \mathcal{P}$ in an interesting class of problems. In particular, we consider cases where \mathcal{H} is an RKHS over a space \mathcal{T} or, more generally, the local parameter space $\dot{\mathcal{H}}_P$ is an RKHS over \mathcal{T} . Denote the feature map of $\dot{\mathcal{H}}_P$ by $t \mapsto K_t$. For $P \in \mathcal{P}$, define $\tilde{\phi}_P : \mathcal{Z} \rightarrow \mathcal{H}$ as follows for each $t \in \mathcal{T}$:

$$\tilde{\phi}_P(z)(t) = \dot{\nu}_P^*(K_t)(z) \quad P\text{-a.s. } z. \quad (5)$$

The following result shows that $\tilde{\phi}_P$ both provides the form of the EIF of ν , when it exists, and also a sufficient condition that can be used to verify this existence.

Theorem 1 (Form of the efficient influence function in RKHS settings). *Suppose ν is pathwise differentiable at P and $\dot{\mathcal{H}}_P$ is an RKHS. Both of the following implications hold:*

- (i) *If ν has an EIF ϕ_P at P , then $\phi_P = \tilde{\phi}_P$ P -almost surely.*
- (ii) *If $\|\tilde{\phi}_P\|_{L^2(P; \mathcal{H})} < \infty$, then ν has an EIF at P .*

The form of the EIF in (5) naturally generalizes its form in finite-dimensional spaces, where the feature K_t replaces the t -th standard basis element e_t . The proof of (i) is a straightforward extension of results about the Riesz representation of a bounded linear functional to our setting, where $\dot{\nu}_P^*(\cdot)(z)$ is only known to be bounded and linear P -almost surely (cf. Lemma 10 of Berlinet and Thomas-Agnan, 2011). The proof of (ii) is more subtle, and involves showing that, when $\|\tilde{\phi}_P\|_{L^2(P; \mathcal{H})} < \infty$, any separable version of the efficient influence process must P -a.s. have sample paths $\dot{\nu}_P^*(\cdot)(z)$ that are both bounded and linear. In the remainder of this subsection, we suppose that ν has an EIF ϕ_P at each $P \in \mathcal{P}$.

In Section 4.1, we establish that, under conditions, a cross-fitted variant of the one-step estimator $\widehat{\nu}_n := \nu(\widehat{P}_n) + P_n\phi_n$ is efficient, in the sense that $\|\widehat{\nu}_n - \nu(P_0)\|_{\mathcal{H}}$ is as concentrated about zero as is possible for any estimator satisfying appropriate regularity conditions. Here, we provide two more heuristic arguments as to why the one-step correction should lead to improvements. The first, which applies specifically in cases where \mathcal{H} is an RKHS, is based on the pointwise performance of the one-step estimator. In particular, the fact that norm convergence in an RKHS implies pointwise convergence can be used to show that the pathwise differentiability of $\nu : \mathcal{P} \rightarrow \mathcal{H}$ also implies the pathwise differentiability of $\nu(\cdot)(t) : \mathcal{P} \rightarrow \mathbb{R}$ for each $t \in \mathcal{T}$. Moreover, the EIF of $\nu(\cdot)(t)$ at $P \in \mathcal{P}$ is equal to $z \mapsto \dot{\nu}_P^*(K_t)(z)$, and so the one-step estimator for the real-valued parameter $\nu(P_0)(t)$ is equal to the evaluation of the \mathcal{H} -valued one-step estimator $\widehat{\nu}_n$ at the point t . While this pointwise justification of the one-step estimator $\widehat{\nu}_n$ is informative, it does not in itself explain why $\widehat{\nu}_n$ should be expected to perform well in a norm sense. Indeed, pointwise convergence in an RKHS does not necessarily imply norm convergence. For the same reason, pathwise differentiability of the real-valued parameters $\nu(\cdot)(t)$, $t \in \mathcal{T}$, does not generally imply pathwise differentiability of the RKHS-valued parameter ν .

The second heuristic justification that we provide here provides initial insights into why the one-step estimator $\widehat{\nu}_n$ should perform well in a norm sense. This justification applies regardless of whether $\dot{\mathcal{H}}_P$ is an RKHS. Fix a submodel $\{P_\epsilon : \epsilon \in [0, \delta]\} \in \mathcal{P}(P, \mathcal{P}, s)$. In the appendix, we establish that, under conditions on either the EIF (Lemma S5) or the submodel (Lemma S6), a first-order approximation to the local parameter $\dot{\nu}_P(s)$ is given by $\epsilon^{-1}(P_\epsilon - P)\phi_P$, in the sense that the difference between these quantities converges to zero as $\epsilon \rightarrow 0$. Combining this with (2) and the fact that $P\phi_P = 0$, this yields the approximation $\nu(P_\epsilon) \approx \nu(P) + P_\epsilon\phi_P$, which is valid up to an additive $o(\epsilon)$ remainder term. Letting P_0 play the role of P_ϵ and \widehat{P}_n play the role of P , this suggests that the von Mises approximation

$$\nu(P_0) \approx \nu(\widehat{P}_n) + P_0\phi_n \tag{6}$$

may also be valid up to a term that goes to zero in probability at a reasonable rate. Admittedly, caution is needed when making the leap from the approximation along the fixed quadratic mean differentiable submodel $\{P_\epsilon : \epsilon \in [0, \delta]\}$ to an approximation that involves the random quantity \widehat{P}_n . For a given parameter ν , the sense in which the above approximation holds can be made precise by directly studying the quantity $\nu(P_0) - \nu(\widehat{P}_n) - P_0\phi_n$. In any case, the approximation above is appealing in that it only relies on P_0 through an expectation, which can naturally be approximated by an expectation under the empirical distribution. This, therefore, suggests the one-step estimator $\widehat{\nu}_n := \nu(\widehat{P}_n) + P_n\phi_n$.

As will follow from the upcoming Lemma 7, $\dot{\mathcal{H}}_P$ need not be an RKHS for an EIF to exist. Consequently, when one does, it is natural to wonder whether there is a general expression for its form. The Riesz representation theorem provides an affirmative answer to this question, showing that, when an EIF exists, it is P -a.s. equal to the following convergent sum:

$$\phi_P(z) = \sum_{k=1}^{\infty} \dot{\nu}_P^*(h_k)(z) h_k, \tag{7}$$

where here and throughout we let $(h_k)_{k=1}^{\infty}$ denote an orthonormal basis of \mathcal{H} . If $\dot{\mathcal{H}}_P$ is an RKHS, evaluating the expression for the EIF in (5) is typically easier than computing (or approximating) the infinite sum in (7). However, in non-RKHS settings, (7) is useful both as an explicit expression for the EIF, if it exists, and as a basis for generalizing the one-step estimator to settings where it does not.

2.4 Regularized one-step estimation when there is no efficient influence function

We now introduce a generalization of the one-step estimator that can be employed regardless of whether an EIF exists. This estimator is a type of series estimator (Schwartz, 1967; Chen, 2007) based on the Riesz representation of a regularized form of the efficient influence operator. This regularized form is motivated by the fact that, when ν has an EIF ϕ_P , it is P -a.s. true that $\dot{\nu}_P^*(\cdot)(z) : h \mapsto \sum_{k=1}^{\infty} \langle h, h_k \rangle_{\mathcal{H}} \dot{\nu}_P^*(h_k)(z)$. The regularized form is designed to ensure that the terms in this sum must decay as k grows sufficiently large. For a square summable $[0, 1]$ -valued sequence $\beta := (\beta_k)_{k=1}^{\infty}$, the β -regularized efficient influence operator is given by $r_P^{\beta}(h)(z) := \sum_{k=1}^{\infty} \beta_k \langle h, h_k \rangle_{\mathcal{H}} \dot{\nu}_P^*(h_k)(z)$. We now show that $r_P^{\beta}(\cdot)(z)$ is always P -a.s. bounded and linear, and we also provide an explicit expression for its Riesz representation. In what follows we let $\ell_*^2 := \ell^2 \cap [0, 1]^{\mathbb{N}}$.

Lemma 1 (β -regularized EIF based on β -regularized efficient influence operator). *If ν is pathwise differentiable at P and $\beta \in \ell_*^2$, then $r_P^{\beta}(\cdot)(z) : \mathcal{H} \rightarrow \mathbb{R}$ is a bounded linear functional on a P -probability one set \mathcal{Z}^{β} with Riesz representation*

$$\phi_P^{\beta}(z) := \sum_{k=1}^{\infty} \beta_k \dot{\nu}_P^*(h_k)(z) h_k.$$

Moreover, $\sigma_P(\beta) := \|\phi_P^{\beta}\|_{L^2(P; \mathcal{H})} = [\sum_{k=1}^{\infty} \beta_k^2 P \dot{\nu}_P(h_k)^2]^{1/2} \leq \|\dot{\nu}_P^*\|_{\text{op}} \|\beta\|_{\ell^2} < \infty$.

When ν has an EIF, ϕ_P^{β} is similar to the expression for it given in (7), but with the k -th term damped by the multiplier $\beta_k \in [0, 1]$. Given this similarity, we call ϕ_P^{β} the β -regularized EIF of ν at P . The corresponding β -regularized one-step estimator is $\hat{\nu}_n^{\beta} := \nu(\hat{P}_n) + P_n \phi_n^{\beta}$.

We now provide a heuristic argument that is similar to one used in the previous subsection for justifying the (non-regularized) one-step estimator, but adapted to account for the regularization bias that arises from using a β -regularized EIF. Fix a submodel $\{P_{\epsilon} : \epsilon\} \in \mathcal{P}(P, \mathcal{P}, s)$. In Lemma S7 in the appendix, we show that, under a regularity condition on the submodel,

$$\begin{aligned} & \left\| \nu(P) - \nu(P_{\epsilon}) + P_{\epsilon} \phi_P^{\beta} - \sum_{k=1}^{\infty} (1 - \beta_k) \langle \nu(P) - \nu(P_{\epsilon}), h_k \rangle_{\mathcal{H}} h_k \right\|_{\mathcal{H}} \\ &= \left(1 + \|\phi_P^{\beta}\|_{L^2(P; \mathcal{H})} \right) \cdot o(\epsilon), \end{aligned} \tag{8}$$

where the $o(\epsilon)$ term does not depend on the choice of β . Similarly to how we did when deriving (6), we let P_0 play the role of P_{ϵ} and \hat{P}_n play the role of P . Recalling that $\hat{\nu}_n^{\beta} := \nu(\hat{P}_n) + P_n \phi_n^{\beta}$ then yields that

$$\hat{\nu}_n^{\beta} - \nu(P_0) \approx (P_n - P_0) \phi_n^{\beta} + \mathcal{B}_{\hat{P}_n}^{\beta}, \tag{9}$$

where, for $P' \in \mathcal{P}$, we let $\mathcal{B}_{P'}^{\beta} := \sum_{k=1}^{\infty} (1 - \beta_k) \langle \nu(P') - \nu(P_0), h_k \rangle_{\mathcal{H}} h_k$. Our formal study of the regularized one-step estimator in Section 5.1 builds on the above. Informally speaking, that study will show that the latter term above plays the role of a regularization bias term that decays as β grows entrywise to $(1, 1, 1, \dots)$ under conditions, and the leading term plays the role of a variance term whose magnitude typically grows with that of β . Hence, a bias-variance tradeoff must be considered when selecting a value for the tuning parameter β . In Section 5.3, we describe a cross-validation strategy for making this selection. There, we also discuss the selection of the basis $(h_k)_{k=1}^{\infty}$.

2.5 Cross-fitted (regularized) one-step estimation

So far, the estimators we have defined have assumed the availability of an iid sample that is independent of Z_1, Z_2, \dots, Z_n that can be used to obtain the estimate \hat{P}_n of P_0 . We now describe how cross-fitting (Schick, 1986; Klaassen, 1987; Zheng and Laan, 2011; Chernozhukov et al., 2018) can be used to avoid the need for this independent sample. For simplicity, we focus on the case of 2-fold cross-fitting and suppose that the sample size is an even number. The generalizations to k -fold cross-fitting ($k \geq 2$) and to the case where n is not divisible by k are straightforward and so are omitted. Let $\hat{P}_n^1 \in \mathcal{P}$ denote an estimate of P_0 based on $\{Z_i\}_{i=1}^{n/2}$ and let P_n^1 denote the empirical distribution of the remainder of the sample $\{Z_i\}_{i=n/2+1}^n$. Define \hat{P}_n^2 and P_n^2 similarly, but with the roles of the two subsamples reversed. We note that, in a slight abuse of notation, P_n^j denotes an empirical distribution derived from $n/2$ observations rather than a j -fold product distribution derived from j independent draws from the empirical distribution P_n of the full sample $\{Z_i\}_{i=1}^n$. Cross-fitting enables the use of arbitrary estimation strategies when constructing \hat{P}_n^j , $j \in \{1, 2\}$, including those based on machine learning techniques.

We now present the form of our cross-fitted estimators. From a notational standpoint, these estimators will be denoted by replacing the hat accents used to denote the sample-splitting estimators in Sections 2.3 and 2.4 by bar accents — for example, the cross-fitted one-step estimator will be denoted by $\bar{\nu}_n$ rather than $\hat{\nu}_n$. This cross-fitted one-step estimator takes the form $\bar{\nu}_n := \frac{1}{2} \sum_{j=1}^2 [\nu(\hat{P}_n^j) + P_n^j \phi_n^j]$, where $\phi_n^j := \phi_{\hat{P}_n^j}$. Let $\phi_n^{j,\beta} := \phi_{\hat{P}_n^j}^\beta$. The cross-fitted β -regularized one-step estimator takes the form $\bar{\nu}_n^\beta := \frac{1}{2} \sum_{j=1}^2 [\nu(\hat{P}_n^j) + P_n^j \phi_n^{j,\beta}]$.

3 Examples of pathwise differentiable parameters

In this section, we present examples of pathwise differentiable Hilbert-valued parameters and the forms of their efficient influence operators and, where applicable, EIFs. From these objects, cross-fitted (regularized) one-step estimators can be derived using the formulas at the end of the previous section. We study the performance of these estimators in Section 6.

In the main text, we focus on parameters that have recently become objects of interest to the causal inference community. Two of these examples (Examples 1b and 3) consider cases where the action space is an RKHS, and two (Examples 1a and 2) consider cases where the action space is an L^2 space, and is therefore not an RKHS. In Appendix A, we show that four more well-studied Hilbert-valued parameters are also pathwise differentiable. In particular, we show that regression functions, square-root density functions, and conditional average treatment effect functions are pathwise differentiable when viewed as elements of appropriate L^2 spaces, and we also show that a kernel mean embedding of a distribution (Gretton et al., 2012) is pathwise differentiable when viewed as an element of an RKHS. We are not aware of any reference establishing the pathwise differentiability of any of the eight Hilbert-valued parameters that we consider in this work.

All derivations for our examples are deferred to Appendix B. For most of these examples, our derivations make use of the following lemma, which we prove in Appendix C. In this lemma, $H(P, P') := [\int (\sqrt{dP} - \sqrt{dP'})^2]^{1/2}$ denotes the Hellinger distance.

Lemma 2 (Sufficient condition for pathwise differentiability). $\nu : \mathcal{P} \rightarrow \mathcal{H}$ is pathwise differentiable at P with local parameter $\nu_P = \eta_P$ if both of the following hold:

- (i) $\eta_P : \dot{\mathcal{P}}_P \rightarrow \mathcal{H}$ is bounded and linear and there exists a set of scores $\mathcal{S}(P)$ whose $L^2(P)$ -closure is equal to $\dot{\mathcal{P}}_P$ such that, for all $s \in \mathcal{S}(P)$, there is at least one submodel $\{P_\epsilon : \epsilon\} \in \mathcal{P}(P, \mathcal{P}, s)$ for which $\|\nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s)\|_{\mathcal{H}} = o(\epsilon)$; and

(ii) ν is locally Lipschitz at P in the sense that there exist $(c, \delta) \in (0, \infty)^2$ such that

$$\|\nu(P_1) - \nu(P_2)\|_{\mathcal{H}} \leq cH(P_1, P_2) \quad \text{for all } P_1, P_2 \in B_{\delta}(P), \quad (10)$$

where $B_{\delta}(P)$ consists of all $P' \in \mathcal{P}$ for which $H(P, P') \leq \delta$.

This lemma is most useful when the set $\mathcal{S}(P)$ and corresponding submodels in $\mathcal{P}(P, \mathcal{P}, s)$ in (i) can be chosen to make establishing (2) for those submodels simple. For example, in a locally nonparametric model, we may take $\mathcal{S}(P)$ to be the set of bounded, P -mean zero functions, and we may take the chosen submodel in $\mathcal{P}(P, \mathcal{P}, s)$ to be such that, for all $\epsilon \in [0, 1/\text{ess sup}_z |s(z)|]$, $\frac{dP_{\epsilon}}{dP} = 1 + \epsilon s$, where the essential supremum is under P .

All of the examples presented in the main text are motivated by questions arising in causal inference. The data structure is common across them, with $Z = (X, A, Y) \sim P$, where X is a vector of covariates with support on \mathcal{X} , A is a treatment with support on either $\{0, 1\}$ or \mathbb{R} , and Y is an outcome with support on \mathcal{Y} . In each example, we suppose that \mathcal{P} is locally nonparametric. We further suppose, for simplicity, that all pairs of distributions in \mathcal{P} are mutually absolutely continuous. For a given distribution P , we let $P_{Y|A,X}$ denote the conditional distribution of Y given (A, X) and P_X denote the marginal distribution of X . We let $g_P(\cdot | x)$ denote the conditional probability mass function of A given $X = x$ under P , when A is binary, and the density of A given $X = x$ under P , when A is continuous. For $s \in \dot{\mathcal{P}}_P$, we let $s_X(x) := E_P[s(Z) | X = x]$ and $s_{Y|A,X}(y | a, x) := s(z) - E_P[s(Z) | A = a, X = x]$.

Example 1a (Counterfactual density function). Suppose that the treatment A is binary and the goal is to estimate the density function of the counterfactual outcome in a setting where everyone receives treatment $A = 1$. This density can offer a more nuanced measure of causal effects than can a more commonly studied counterfactual mean outcome (Kennedy et al., 2021). Suppose that there is a σ -finite measure λ_Y such that, for all $P \in \mathcal{P}$, there is a regular conditional probability $P_{Y|A,X}$ such that $P_{Y|A,X}(\cdot | a, x) \ll \lambda_Y$ for P -almost all $(a, x) \in \{0, 1\} \times \mathcal{X}$. Define the propensity to receive treatment a as $g_P(a | x) := P(A = a | X = x)$ and let $p_{Y|A,X}(\cdot | 1, x)$ denote the conditional density of Y given $(A, X) = (1, x)$. The parameter of interest $\nu : \mathcal{P} \rightarrow L^2(\lambda_Y)$ takes the form

$$\nu(P)(y) = \int p_{Y|A,X}(y | 1, x) P_X(dx). \quad (11)$$

Under typical causal assumptions, ν corresponds to the density of the counterfactual outcome that would be seen under treatment $A = 1$. We require that \mathcal{P} satisfy the following conditions:

$$\inf_{P \in \mathcal{P}} \text{ess inf}_{x \in \mathcal{X}} g_P(1 | x) > 0 \quad \text{and} \quad \sup_{P \in \mathcal{P}} \text{ess sup}_{(x,y) \in \mathcal{X} \times \mathcal{Y}} p_{Y|A,X}(y | 1, x) < \infty, \quad (12)$$

where the essential infimum is under P_X and the essential supremum is under $P_X \times \lambda_Y$. The first inequality, referred to as a strong positivity assumption, holds when the propensity to receive treatment 1 is not vanishingly small. The second holds when the conditional density of Y given $(A, X) = (1, x)$ is uniformly bounded across all distributions in the model. While in principle the second condition could be weakened, this condition appears to be sufficiently general to capture many statistical models of interest.

The local parameter takes the form

$$\dot{\nu}_P(s)(y) = \int \{s_{Y|A,X}(y | 1, x) + s_X(x)\} p_{Y|A,X}(y | 1, x) P_X(dx), \quad (13)$$

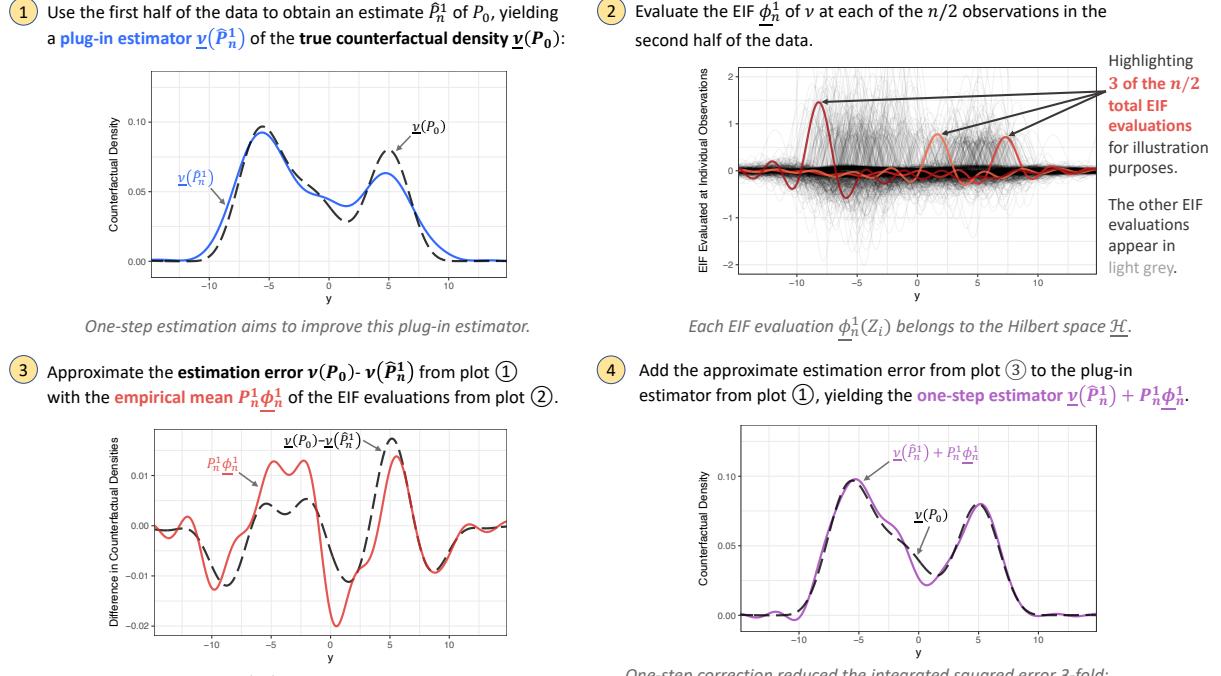


Figure 1: Illustration of how a sample-splitting one-step estimator is constructed in Example 1b. The cross-fitted estimator averages two such estimators, one obtained as above and the other with the roles of the two halves of the data reversed. Details of the data-generating process, along with a Monte Carlo assessment of the performance of a cross-fitted one-step estimator, are given in Appendix I.1.

and the efficient influence operator takes the form

$$\begin{aligned} \dot{\nu}_P^*(h)(y, a, x) &= \frac{1\{a = 1\}}{g_P(a | x)} \{h(y) - E_P [h(Y) | A = a, X = x]\} \\ &+ \left(E_P [h(Y) | A = 1, X = x] - \int E_P [h(Y) | A = 1, X = x'] P_X(dx') \right). \end{aligned} \quad (14)$$

Unless λ_Y is a discrete measure, $\dot{\nu}_P^*$ will not generally be a bounded operator. This can be shown to follow from the facts that point evaluation is not continuous in L^2 spaces (page 8 of [Berlinet and Thomas-Agnan, 2011](#)) and $\dot{\nu}_P^*(h)(y, a, x)$ depends on the evaluation of h at y .

Example 1b (Bandlimited counterfactual density function). The setting is the same as in Example 1a, except that Y must be real-valued and continuous, λ_Y must be the Lebesgue measure, and, for fixed $b > 0$, the target of inference is the following transformation of the counterfactual density $\nu(P)$ that was defined in Example 1a:

$$\underline{v}(P)(y) := \int_{-\infty}^{\infty} K_y(\tilde{y}) \nu(P)(\tilde{y}) \lambda_Y(d\tilde{y}), \quad (15)$$

where $K_y(\tilde{y}) := \{\sin[b(\tilde{y} - y)]\}/[\pi(\tilde{y} - y)]$. This estimand corresponds to a bandlimiting of the counterfactual density function. In particular, letting \mathcal{F} and \mathcal{F}^{-1} denote the Fourier transform and inverse Fourier transform and fixing $b > 0$, the b -bandlimiting of a Lebesgue square integrable function $f : \mathbb{R} \rightarrow \mathbb{R}$ is the function $\mathcal{B}(f) : \mathbb{R} \rightarrow \mathbb{R}$ given by

$$\mathcal{B}(f) := \mathcal{F}^{-1}(1_{[-b, b]} \cdot \mathcal{F}(f))(y) = \int_{-\infty}^{\infty} K_y(\tilde{y}) f(\tilde{y}) \lambda_Y(d\tilde{y}), \quad (16)$$

where $1_{[-b,b]} \cdot \mathcal{F}(f)$ represents the function $\xi \mapsto 1_{[-b,b]}(\xi) \cdot \mathcal{F}(f)(\xi)$ and the latter equality holds by the convolution theorem. The estimand is equal to $\underline{\nu}(P) := \mathcal{B}(\nu(P))$. Lemma S1 in Appendix B.2 shows that $\mathcal{B}(\nu(P))$ corresponds to an $L^2(\lambda_Y)$ projection of $\nu(P)$ onto

$$\underline{\mathcal{H}} := \{h \in C(\mathbb{R}) \mid \text{support}[\mathcal{F}(h)] \subseteq [-b, b]\},$$

where $C(\mathbb{R})$ denotes the set of continuous, Lebesgue square integrable functions. A little care is needed to make this result precise since $L^2(\lambda_Y)$ is a space of equivalence classes of functions whereas $\underline{\mathcal{H}}$ is a space of functions (see the statement of Lemma S1 for details). The space $\underline{\mathcal{H}}$ is an RKHS when equipped with the $L^2(\lambda_Y)$ inner product $\langle h, \tilde{h} \rangle_{\underline{\mathcal{H}}} = \int h(y) \tilde{h}(y) \lambda_Y(dy)$ (Yao, 1967). The kernel function in this RKHS is given by $(y, \tilde{y}) \mapsto \underline{K}_y(\tilde{y})$. This RKHS is a smoothness class consisting of all square integrable functions that have an analytic continuation to the complex plane that satisfies an exponential growth restriction (Theorem 19.3 and page 372 of Rudin, 1987). When a (non-counterfactual) density function belongs to $\underline{\mathcal{H}}$, a particular kernel density estimator has been shown to attain mean integrated squared error (MISE) that decays at an n^{-1} rate (Ibragimov and Khas'minskii, 1983; Agarwal et al., 2015). Our setting differs from earlier ones in that (i) we focus on a counterfactual density, and (ii) we define our estimand as a nonparametric projection onto the space of b -bandlimited functions, rather than requiring that our estimand belong to this class. Naturally, when the counterfactual density $\nu(P)$ already belongs to $\underline{\mathcal{H}}$, $\underline{\nu}(P)$ will equal this density and so our approach will yield estimators of it.

In the appendix, we rely heavily on the calculations performed in Example 1a when showing that $\underline{\nu} : \mathcal{P} \rightarrow \underline{\mathcal{H}}$ is pathwise differentiable. We show that the local parameter $\dot{\underline{\nu}}_P : \dot{\mathcal{P}}_P \rightarrow \underline{\mathcal{H}}$ of $\underline{\nu}$ at P is closely related to that of ν . In particular, $\dot{\underline{\nu}}_P(s) = \mathcal{B}(\dot{\nu}_P(s))$, where $\dot{\nu}_P$ is as defined in Example 1a. Letting $[h]$ denote the equivalence class of functions that are equal to some $h \in \underline{\mathcal{H}}$ Lebesgue-almost everywhere, the efficient influence operator $\dot{\nu}_P^* : \underline{\mathcal{H}} \rightarrow \dot{\mathcal{P}}_P$ takes the form $\dot{\underline{\nu}}_P^*(h) = \dot{\nu}_P^*([h])$, where $\dot{\nu}_P^*$ is as defined in Example 1a. Since $\underline{\mathcal{H}}$ is an RKHS, we can also look to define the EIF $\underline{\phi}_P$ of $\underline{\nu}$ at P . In particular, (12) can be used to show that the function $\underline{\phi}_P(z) : y \mapsto \underline{\nu}_P^*(\underline{K}_y)(z)$ belongs to $L^2(P; \underline{\mathcal{H}})$, and so $\underline{\phi}_P$ is indeed the EIF of $\underline{\nu}$ at P . See (S5) in the appendix for a more explicit expression for the EIF $\underline{\phi}_P$.

Figure 1 shows how a one-step estimator can improve an initial estimator in this example.

Example 2 (Counterfactual mean outcome under a continuous treatment). In the previous example, the treatment was considered to be binary. In this example, we take A to be a continuous treatment taking values in $\mathcal{A} = [0, 1]$, such as a dosage, duration, or frequency of intervention. Denote the marginal distribution of A by P_A and suppose that the Lebesgue measure λ_A dominates the conditional distribution $P_{A|X}(\cdot \mid x)$ of $A \mid X = x$ under P for P -almost all x . Let $g_P(\cdot \mid x)$ denote the conditional density of A given that $X = x$. The target of estimation is $\nu : \mathcal{P} \rightarrow L^2(\lambda_A)$, where $\nu(P)(a) = \int E_P[Y \mid A = a, X = x] P_X(dx)$. Under causal conditions, $\nu(P)(a)$ corresponds to the mean outcome under a continuous treatment (Díaz and van der Laan, 2013). We suppose the strong positivity assumption that $\inf_{P \in \mathcal{P}} \text{ess inf}_{(a,x)} g_P(a \mid x) > 0$, where the essential infimum is under $\lambda_A \times P_X$. We further suppose that Y has a bounded conditional second moment, in the sense that $\sup_{P \in \mathcal{P}} \text{ess sup}_{(a,x)} E_P[Y^2 \mid A = a, X = x] < \infty$, where the essential supremum is under $\lambda_A \times P_X$.

The local parameter takes the form

$$\begin{aligned} \dot{\nu}_P(s)(a) &= \iint \{y - \mu_P(a, x)\} s_{Y|A,X}(y \mid a, x) P_{Y|A,X}(dy \mid a, x) P_X(dx) \\ &\quad + \int [\mu_P(a, x) - \nu(P)(a)] s_X(x) P_X(dx), \end{aligned} \tag{17}$$

where $\mu_P(a, x) = E_P[Y \mid A = a, X = x]$. The efficient influence operator takes the form

$$\dot{\nu}_P^*(h)(y, a, x) = \frac{y - \mu_P(a, x)}{g_P(a \mid x)} h(a) + \int [\mu_P(a', x) - \nu(P)(a')] h(a') \lambda_A(da'). \quad (18)$$

Similarly to Example 1a, $\dot{\nu}_P^*$ is not generally a bounded operator.

Example 3 (Counterfactual kernel mean embedding). Let A be a binary treatment and suppose the strong positivity assumption that $\inf_{P \in \mathcal{P}} \text{ess inf}_{x \in \mathcal{X}} g_P(1 \mid x) > 0$, where the essential infimum is under P_X . Let $\kappa : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ be a bounded, symmetric, positive definite function, \mathcal{H} be the RKHS associated with the kernel κ , and $K_y := \kappa(y, \cdot)$ be the associated feature map. The counterfactual kernel mean embedding (Muandet et al., 2021; Fawkes et al., 2022) is a parameter $\nu : \mathcal{P} \rightarrow \mathcal{H}$ such that

$$\nu(P) = \int E_P[K_Y \mid A = 1, X = x] P_X(dx).$$

Under standard causal conditions (Robins, 1986), $\nu(P)$ can be shown to be equal to the kernel mean embedding (Gretton et al., 2012) of the distribution of a counterfactual outcome in a world where treatment 1 was given to everyone. We suppose that the strong positivity assumption in (12) holds.

The local parameter takes the form

$$\dot{\nu}_P(s) = \iint K_y [s_{Y|A,X}(y \mid 1, x) + s_X(x)] P_{Y|A,X}(dy \mid 1, x) P_X(dx). \quad (19)$$

The efficient influence operator takes the form

$$\begin{aligned} \dot{\nu}_P^*(h)(y, a, x) &= \frac{a}{g_P(1 \mid x)} \{h(y) - E_P[h(Y) \mid A = a, X = x]\} \\ &\quad + E_P[h(Y) \mid A = 1, X = x] - \int E_P[h(Y) \mid A = 1, X = x'] P_X(dx'), \end{aligned}$$

and the EIF is P -Bochner square integrable and takes the form

$$\phi_P(y, a, x) = \frac{a}{g_P(1 \mid x)} \{K_y - \mu_P^K(x)\} + \mu_P^K(x) - \nu(P), \quad (20)$$

where $\mu_P^K : \mathcal{X} \rightarrow \mathcal{H}$ is defined as $\mu_P^K(x) := E_P[K_Y \mid A = 1, X = x]$.

4 Performance guarantees and inference when there is an EIF

4.1 Performance guarantees for one-step estimators

In this section, we provide conditions under which a cross-fitted one-step estimator $\bar{\nu}_n$ is both asymptotically linear and efficient. These conditions concern the terms arising in the following decomposition:

$$\begin{aligned} \|\bar{\nu}_n - \nu(P_0) - P_n \phi_0\|_{\mathcal{H}} &= \frac{1}{2} \left\| \sum_{j=1}^2 \left[\nu(\hat{P}_n^j) + P_0 \phi_n^j - \nu(P_0) \right] + \sum_{j=1}^2 (P_n^j - P_0)(\phi_n^j - \phi_0) \right\|_{\mathcal{H}} \\ &\leq \max_j \|\mathcal{R}_n^j\|_{\mathcal{H}} + \max_j \|\mathcal{D}_n^j\|_{\mathcal{H}}, \end{aligned} \quad (21)$$

where $\mathcal{R}_n^j := \nu(\hat{P}_n^j) + P_0 \phi_n^j - \nu(P_0)$ and $\mathcal{D}_n^j := (P_n^j - P_0)(\phi_n^j - \phi_0)$. We call \mathcal{R}_n^j the remainder terms and \mathcal{D}_n^j the drift terms, $j \in \{1, 2\}$. Asymptotic linearity, as defined in (4), holds whenever both of these quantities are negligible, in the sense that they are $o_p(n^{-1/2})$. The remainder terms \mathcal{R}_n^j , $j \in \{1, 2\}$, quantify the error in the approximation (6) across the two splits of the sample. As described heuristically

in the text surrounding (6), it is reasonable to expect that this remainder term will be negligible under appropriate conditions. The following result provides a reasonable condition under which the drift terms will be negligible.

Lemma 3 (Sufficient condition for negligible drift terms). *Suppose that ν is pathwise differentiable at P_0 with EIF ϕ_0 . For each $j \in \{1, 2\}$, $\|\phi_n^j - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$ implies that $\|\mathcal{D}_n^j\|_{\mathcal{H}} = o_p(n^{-1/2})$.*

In the appendix, this lemma is proved via a conditioning argument that makes use of Chebyshev's inequality for Hilbert-valued random variables (Grenander, 1963) and the dominated convergence theorem. In the next section, we provide lower-level conditions under which the remainder and drift terms will be small in the contexts of Examples 1b and 3.

In the process of showing that $\bar{\nu}_n$ is asymptotically linear, we also show that it is regular. An estimator sequence $(\tilde{\nu}_n)_{n=1}^\infty$ is said to be regular at $P_0 \in \mathcal{P}$ if and only if there is a tight \mathcal{H} -valued random variable $\tilde{\mathbb{H}}$ such that, for every score in the tangent set, quadratic mean differentiable submodel $\{P_\epsilon : \epsilon\} \in \mathcal{P}(P_0, \mathcal{P}, s)$, and every sequence $\epsilon_n = O(n^{-1/2})$, $\sqrt{n}[\tilde{\nu}_n - \nu(P_{\epsilon_n})]$ converges weakly to $\tilde{\mathbb{H}}$ under iid sampling of n observations from P_{ϵ_n} . We say that an estimator $\tilde{\nu}_n$ is regular when the implied estimator sequence $(\tilde{\nu}_n)_{n=1}^\infty$ is clear from context. In the upcoming theorem, we write ' \rightsquigarrow ' to denote weak convergence in \mathcal{H} .

Theorem 2 (Asymptotic linearity and weak convergence of a one-step estimator). *Suppose that ν is pathwise differentiable at P_0 with EIF $\phi_0 \in L^2(P_0; \mathcal{H})$ and, for $j \in \{1, 2\}$, $\mathcal{R}_n^j = o_p(n^{-1/2})$ and $\mathcal{D}_n^j = o_p(n^{-1/2})$. Under these conditions, (4) holds, $\bar{\nu}_n$ is regular, and*

$$n^{1/2} [\bar{\nu}_n - \nu(P_0)] \rightsquigarrow \mathbb{H}, \quad (22)$$

where \mathbb{H} is a tight \mathcal{H} -valued Gaussian random variable that is such that, for each $h \in \mathcal{H}$, the marginal distribution $\langle \mathbb{H}, h \rangle_{\mathcal{H}}$ is $N(0, E_0[\langle \phi_0(Z), h \rangle_{\mathcal{H}}^2])$.

The convolution theorem for Banach-valued estimators can be used to characterize one sense in which $\bar{\nu}_n$ is an efficient estimator of $\nu(P_0)$ (see Theorem 3.11.2 and Lemma 3.11.4 of van der Vaart and Wellner, 1996, for a convenient version). Rather than present the theorem in full generality, here we present a particularly interpretable consequence of it. Specifically, under the conditions of Theorem 2, that theorem shows that $\bar{\nu}_n$ is optimally concentrated about $\nu(P_0)$ in the sense that, for any regular estimator sequence $(\tilde{\nu}_n)_{n=1}^\infty$ and $c \geq 0$,

$$\lim_{n \rightarrow \infty} P_0^n \{n\|\bar{\nu}_n - \nu(P_0)\|_{\mathcal{H}}^2 > c\} \leq \lim_{n \rightarrow \infty} P_0^n \{n\|\tilde{\nu}_n - \nu(P_0)\|_{\mathcal{H}}^2 > c\}. \quad (23)$$

The above describes a sense in which $\bar{\nu}_n$ is optimal among all regular estimators of $\nu(P_0)$. Under the conditions of Theorem 2, $\bar{\nu}_n$ can also be shown to outperform all estimators, including non-regular ones, in a local asymptotic minimax sense — see Theorem 3.11.5 in van der Vaart and Wellner (1996) for details.

4.2 Construction of confidence sets based on one-step estimators

As we will now show, the weak convergence result in Theorem 2 can be used to facilitate the construction of $(1 - \alpha)$ -level confidence sets for the Hilbert-valued parameter $\nu(P_0)$, where $\alpha \in (0, 1)$ is some fixed constant. Our proposed confidence set is built based upon a quadratic form $w(\cdot; \Omega) : h \mapsto \langle \Omega(h), h \rangle_{\mathcal{H}}$ that is parameterized by a standardization operator Ω that belongs to the set \mathcal{O} of continuous, self-adjoint, positive definite linear operators mapping from \mathcal{H} to \mathcal{H} . In particular, letting $\zeta \geq 0$ be a specified threshold and $\Omega_n \in \mathcal{O}$ an estimator of a some possibly- P_0 -dependent operator $\Omega_0 \in \mathcal{O}$, our confidence set

will take the form

$$\mathcal{C}_n(\zeta) := \{h \in \mathcal{H} : w(\bar{\nu}_n - h; \Omega_n) \leq \zeta/n\}. \quad (24)$$

We will see that, for an appropriate choice of ζ and provided $\|\Omega_n - \Omega_0\|_{\text{op}} = o_p(1)$, the continuous mapping theorem justifies the asymptotic validity of this confidence set.

Before presenting that result, we discuss two interesting choices of Ω_0 . The first is the identity function. This choice yields a spherical confidence set that consists of all h belonging to an \mathcal{H} -ball centered at $\bar{\nu}_n$ of radius $(\zeta/n)^{1/2}$. Since the form of Ω_0 does not rely on P_0 in this case, Ω_n can be taken to be equal to Ω_0 . The second is proportional to the inverse of a regularized form of the covariance operator of \mathbb{H} , which will yield a Wald-type confidence set for $\nu(P_0)$ that has an elliptical shape. Compared to spherical confidence sets, Wald-type confidence sets have the benefit of being narrower in the direction of unit vectors h where estimation is easier, in the sense that the variance of $\langle \mathbb{H}, h \rangle_{\mathcal{H}}$ is smaller. Regularization is needed when defining a Wald-type confidence set because the covariance operator of \mathbb{H} will not generally be invertible and, even if it is, this inverse will not be continuous when $\|\phi_0\|_{L^2(P_0; \mathcal{H})} < \infty$ unless \mathcal{H} is finite-dimensional. A simple example of a regularized covariance operator Ω_0 , and an operator-norm consistent estimator thereof, is given in Appendix D. Many other types of regularization are also possible (Tikhonov et al., 1995).

The following result establishes the asymptotic validity of the confidence set in (24) based on a threshold $\hat{\zeta}_n$ that is measurable with respect to the σ -field generated by the iid sample Z_1, \dots, Z_n from P_0 . This threshold is an estimate of the $(1 - \alpha)$ -quantile $\zeta_{1-\alpha}$ of $w(\mathbb{H}; \Omega_0)$.

Theorem 3 (Asymptotically valid confidence set). *Suppose the conditions of Theorem 2 hold. Further suppose that $\|\phi_0\|_{L^2(P_0; \mathcal{H})} > 0$, $\Omega_n \in \mathcal{O}$, $\Omega_0 \in \mathcal{O}$, and $\|\Omega_n - \Omega_0\|_{\text{op}} = o_p(1)$.*

- (i) if $\hat{\zeta}_n \rightarrow \zeta_{1-\alpha}$ in probability, then $\lim_{n \rightarrow \infty} P_0^n \{\nu(P_0) \in \mathcal{C}_n(\hat{\zeta}_n)\} = 1 - \alpha$.
- (ii) if $\hat{\zeta}_n$ is an asymptotically conservative estimator of $\zeta_{1-\alpha}$, in the sense that $P_0^n \{\hat{\zeta}_n \geq \zeta_{1-\alpha} - \delta\} \xrightarrow{n \rightarrow \infty} 1$ for all $\delta > 0$, then $\liminf_{n \rightarrow \infty} P_0^n \{\nu(P_0) \in \mathcal{C}_n(\hat{\zeta}_n)\} \geq 1 - \alpha$.

The proof of this result is similar in structure to those used to establish the validity of Wald-type confidence sets for finite-dimensional parameters. It consists in applying the continuous mapping theorem and Slutsky's lemma to show that $n \cdot w[\bar{\nu}_n - \nu(P_0); \Omega_n] \rightsquigarrow w(\mathbb{H}; \Omega_0)$, showing that $w(\mathbb{H}; \Omega_0)$ is a continuous random variable, and finally using that convergence in distribution implies convergence of cumulative distribution functions at continuity points.

A consistent estimator of $\zeta_{1-\alpha}$ can be defined using the bootstrap (Efron, 1979). To define this estimator, we let $Z_1^\sharp, \dots, Z_{n/2}^\sharp \stackrel{\text{iid}}{\sim} P_n^2$ be sampled independently of $Z_{n/2+1}^\sharp, \dots, Z_n^\sharp \stackrel{\text{iid}}{\sim} P_n^1$. We then let $P_n^{2,\sharp}$ be the empirical distribution of $Z_1^\sharp, \dots, Z_{n/2}^\sharp$ and $P_n^{1,\sharp}$ be the empirical distribution of $Z_{n/2+1}^\sharp, \dots, Z_n^\sharp$. We let $\mathbb{H}_n^\sharp := n^{1/2} \sum_{j=1}^2 (P_n^{j,\sharp} - P_n^j) \phi_n^j / 2$. The threshold $\hat{\zeta}_n$ is taken to be equal to the $(1 - \alpha)$ -quantile of $w(\mathbb{H}_n^\sharp, \Omega_n)$, conditionally on the original sample (Z_1, \dots, Z_n) used to define P_n^1 and P_n^2 . In practice this quantile can be well-approximated by selecting m Monte Carlo draws of the sample $Z_1^\sharp, \dots, Z_n^\sharp$ and then returning the empirical $(1 - \alpha)$ -quantile of $w(\mathbb{H}_n^\sharp, \Omega_n)$ over these draws. A computational benefit of this bootstrap procedure is that it does not require refitting the initial estimators \hat{P}_n^j of P_0 ; indeed, the same EIFs $\phi_n^j := \phi_{\hat{P}_n^j}$ are used for each bootstrap replication.

Theorem 4 (Consistent estimation of $\zeta_{1-\alpha}$ via the bootstrap). *Suppose the conditions of Theorem 2 hold. Further suppose that $\Omega_n \in \mathcal{O}$, $\Omega_0 \in \mathcal{O}$, $\|\Omega_n - \Omega_0\|_{\text{op}} = o_p(1)$, and $\max_{j \in \{1, 2\}} \|\phi_n^j - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$. Under these conditions, $\hat{\zeta}_n \rightarrow \zeta_{1-\alpha}$ in probability.*

In brief, the proof of the above consists in showing that \mathbb{H}_n^\sharp is asymptotically equivalent to $\mathbb{H}_{n,0}^\sharp := n^{1/2} \frac{1}{2} \sum_{j=1}^2 (P_n^{j,\sharp} - P_n^j) \phi_0$ in probability, invoking a guarantee from Giné and Zinn (1990) regarding the

weak convergence of the bootstrap for Hilbert-valued sample means and the continuous mapping theorem to show that $w(\mathbb{H}_{n,0}^\sharp; \Omega_0) \rightsquigarrow w(\mathbb{H}; \Omega_0)$ conditionally on $(Z_i)_{i=1}^\infty$ with probability one, and finally applying a Slutsky-type argument to replace the P_0 -dependent quantities $\mathbb{H}_{n,0}^\sharp$ and Ω_0 in $w(\mathbb{H}_{n,0}^\sharp; \Omega_0)$ by \mathbb{H}_n^\sharp and Ω_n , respectively.

When Ω_0 is the identity function, Theorem 1 in [Székely and Bakirov \(2003\)](#) provides a means to derive an alternative estimator of $\zeta_{1-\alpha}$. This estimator does not require the bootstrap, but is asymptotically conservative. See Appendix E for details.

In practice, it will typically be necessary to use numerical techniques to compute the quadratic form $w(\bar{\nu}_n - h; \Omega_n)$ that is used to define our confidence set. We discuss some such approaches in Appendix F.

5 Performance guarantees and inference when there is no EIF

5.1 Performance guarantees for regularized one-step estimation

In this subsection, we provide performance guarantees for the cross-fitted β_n -regularized one-step estimator $\bar{\nu}_n^{\beta_n}$, where, for each n , β_n is an ℓ_*^2 -valued regularization parameter. Before doing so, we acknowledge a minor abuse of notation. We will denote the k^{th} entry of a generic regularization parameter $\beta \in \ell_*^2$ by β_k , which should not be mistaken for the sample-size- n dependent regularization parameter β_n , whose k^{th} entry we will denote by $\beta_{n,k}$. This should not cause confusion, as we always denote sample size by n and a generic index of a vector in ℓ_*^2 by k .

We will show that, under conditions, $\bar{\nu}_n^{\beta_n}$ satisfies the following biased and slower-than- $n^{-1/2}$ -rate asymptotically linear expansion, which formalizes the approximation in (9) for a cross-fitted one-step estimator:

$$\bar{\nu}_n^{\beta_n} - \nu(P_0) = \frac{1}{2} \sum_{j=1}^2 \mathcal{B}_n^{j,\beta_n} + P_n \phi_0^{\beta_n} + O_p \left(\|\beta_n\|_{\ell^2} / n^{1/2} \right). \quad (25)$$

Above $\phi_0^{\beta_n}$ is as defined in Lemma 1 and $\mathcal{B}_n^{j,\beta_n} := \mathcal{B}_{\hat{P}_n^j}^{\beta_n}$ denotes the bias term defined below (9). When ν does not have an EIF, there is generally a tradeoff between the bias term, whose magnitude is smaller when the entries of β_n are closer to 1, and the linear ‘variance’ term $P_n \phi_0^{\beta_n}$, whose magnitude scales as $O_p[\sigma_0(\beta_n)/n^{1/2}]$, where $\sigma_0(\beta_n) = O(\|\beta_n\|_{\ell^2})$ is as defined in Lemma 1. These two terms will typically be of the same order when β_n is selected to minimize the mean-squared error $E_{P_0^n} [\|\bar{\nu}_n^{\beta_n} - \nu(P_0)\|_{\mathcal{H}}^2]$, which makes it so that $\bar{\nu}_n^{\beta_n} - \nu(P_0)$ converges to zero in probability slower than does $n^{-1/2}$. Owing to the bias term, and also to the fact that there is generally not a scaling of $P_n \phi_0^{\beta_n}$ that will converge to a nondegenerate, tight random element when ν does not have an EIF (see Lemma S11 in the appendix), our focus in this subsection will be on deriving rates of convergence for the regularized estimator $\bar{\nu}_n^{\beta_n}$. We provide a means to construct confidence sets for $\nu(P_0)$ in Section 5.2.

To establish (25), we introduce regularized versions of the drift and remainder terms considered in Section 4.1. In particular, for $j \in \{1, 2\}$ and $\beta = (\beta_k)_{k=1}^\infty \in \ell^2$, define the \mathcal{H} -valued random elements $\mathcal{D}_n^{j,\beta} := (P_n^j - P_0)(\phi_n^{j,\beta} - \phi_0^\beta)$ and $\mathcal{R}_n^{j,\beta} := \mathcal{R}_{\hat{P}_n^j}^\beta$, where, for $P \in \mathcal{P}$,

$$\mathcal{R}_P^\beta := \nu(P) - \nu(P_0) + P_0 \phi_P^\beta - \sum_{k=1}^\infty (1 - \beta_k) \langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} h_k. \quad (26)$$

The main result of this subsection is as follows.

Theorem 5 (Rate of convergence of regularized one-step estimator). *Suppose ν is pathwise differentiable at P_0 , $\beta_n \in \ell_*^2$ for each $n \in \mathbb{N}$, and both $\mathcal{R}_n^{j,\beta_n}$ and $\mathcal{D}_n^{j,\beta_n}$ are $O_p[\|\beta_n\|_{\ell^2}/n^{1/2}]$ for $j \in \{1, 2\}$. Under these*

conditions, (25) holds. Moreover, if $\mathcal{B}_n^{j,\beta_n} = O_p(\|\beta_n\|_{\ell^2}/n^{1/2})$ for $j \in \{1, 2\}$, then

$$\bar{\nu}_n^{\beta_n} - \nu(P_0) = O_p\left(\|\beta_n\|_{\ell^2}/n^{1/2}\right). \quad (27)$$

Eq. 27 suggests it is desirable to select β_n as small as possible, while still ensuring that the drift, remainder, and bias terms are all $O_p(\|\beta_n\|_{\ell^2}/n^{1/2})$. Below we provide three general results that can aid in establishing these conditions. The first two provide ways to guarantee the drift and remainder terms are of no larger order than the variance term $P_n \phi_0^{\beta_n}$ in (25), implying that the rate of convergence is determined by the variance and bias terms. The third makes precise our earlier statement that the bias term is smaller when the entries of β_n are closer to 1, thereby enforcing a lower bound on how small β_n can be to ensure that the bias term is of the same order as the variance term and, therefore, (27) holds.

Lemma 4 (Sufficient condition for negligible regularized drift terms). *Suppose that ν is pathwise differentiable at P_0 and $(r_n)_{n=1}^\infty$ is a nonnegative sequence. Fix $j \in \{1, 2\}$ and, for each $n \in \mathbb{N}$, let $\beta_n \in \ell_*^2$. If $\|\phi_n^{j,\beta_n} - \phi_0^{\beta_n}\|_{L^2(P_0; \mathcal{H})} = o_p(r_n)$, then $\|\mathcal{D}_n^{j,\beta_n}\|_{\mathcal{H}} = o_p(r_n/n^{1/2})$.*

By taking $r_n = \|\beta_n\|_{\ell^2}$, the above gives a condition for $\mathcal{D}_n^{j,\beta_n}$ to be $o_p(\|\beta_n\|_{\ell^2}/n^{1/2})$, and therefore $O_p(\|\beta_n\|_{\ell^2}/n^{1/2})$. For the regularized remainder term, the following can be useful.

Lemma 5 (Bound on regularized remainder term). *Fix $j \in \{1, 2\}$ and $\beta \in \ell_*^2$. If ν is pathwise differentiable at $P \in \mathcal{P}$, $P_0 \ll P$, and $\phi_P^\beta \in L^2(P_0; \mathcal{H})$, then*

$$\|\mathcal{R}_P^\beta\|_{\mathcal{H}}^2 = \sum_{k=1}^{\infty} \beta_k^2 \cdot (\mathcal{R}_{P,k})^2 \leq \|\beta\|_{\ell^2}^2 \sup_{k \in \mathbb{N}} (\mathcal{R}_{P,k})^2,$$

where $\mathcal{R}_{P,k} := \langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} + P_0 \dot{\nu}_P^*(h_k)$.

For a given $k \in \mathbb{N}$, $\mathcal{R}_{P,k}$ corresponds to the remainder term in a von Mises expansion of the real-valued parameter $\psi_k : P' \mapsto \langle \nu(P'), h_k \rangle_{\mathcal{H}}$ (von Mises, 1947), where we note that the pathwise differentiability of ν at P implies the pathwise differentiability of ψ_k at P with canonical gradient $\dot{\nu}_P^*(h_k)$.

We now turn to the bias term. For $u \geq 0$, let $\|\cdot\|_u : \mathcal{H} \rightarrow [0, +\infty]$ denote the norm defined by $\|h\|_u^2 := \sum_{k=1}^{\infty} k^{2u} \langle h, h_k \rangle_{\mathcal{H}}^2$, where the dependence of $\|\cdot\|_u$ on the basis $(h_k)_{k=1}^\infty$ used to construct the regularized one-step estimator is suppressed in the notation.

Lemma 6 (Bound on bias term). *For any $u \geq 0$ and $\beta \in \ell_*^2$, $\|\mathcal{B}_P^\beta\|_{\mathcal{H}} \leq \|\nu(P) - \nu(P_0)\|_u \sup_{k \in \mathbb{N}} (1 - \beta_k)/k^u$. If there exists $K \in \mathbb{N}$ such that $\beta_k = 1$ for all $k \leq K$ and $\beta_k = 0$ for all $k > K$, then this implies that*

$$\|\mathcal{B}_P^\beta\|_{\mathcal{H}} \leq (K+1)^{-u} \|\nu(P) - \nu(P_0)\|_u \leq 2(K+1)^{-u} \sup_{P' \in \mathcal{P}} \|\nu(P')\|_u. \quad (28)$$

Naturally, the upper bounds are only informative if the evaluations of $\|\cdot\|_u$ upon which they rely are finite. Conditions for the finiteness of this norm have been evaluated in several settings. In particular, $\{h \in \mathcal{H} : \|h\|_u < \infty\}$ corresponds to a periodic Sobolev space when $\mathcal{H} = L^2([0, 1])$ and $(h_k)_{k=1}^\infty$ is the trigonometric basis (Proposition 1.14 of Tsybakov, 2009), Sobolev-Laguerre space when $\mathcal{H} = L^2([0, \infty))$ and $(h_k)_{k=1}^\infty$ consists of the Laguerre functions (Bongioanni and Torrea, 2008), and Sobolev-Hermite space when $\mathcal{H} = L^2(\mathbb{R})$ and $(h_k)_{k=1}^\infty$ consists of the Hermite functions (Bongioanni and Torrea, 2006).

In Section 6, we study the selection of β_n in the context of our examples. We do this by leveraging the bounds from the preceding three lemmas and then deriving the choice of β_n that balances the variance and bias terms.

5.2 Construction of confidence sets

In what follows we fix $\beta \in \ell_*^2$ and define $\Gamma_\beta : \mathcal{H} \rightarrow \mathcal{H}$ as $\Gamma_\beta(h) = \sum_{k=1}^\infty \beta_k \langle h, h_k \rangle_{\mathcal{H}} h_k$. The following is the key observation that we use to construct our confidence set for $\nu(P_0)$.

Lemma 7. *If ν is pathwise differentiable at P , then its transformation $\nu^\beta := \Gamma_\beta \circ \nu$ is pathwise differentiable at P with local parameter $\dot{\nu}_P^\beta := \Gamma_\beta \circ \dot{\nu}_P$ and EIF $\phi_P^\beta \in L^2(P; \mathcal{H})$.*

Since ν^β has an EIF, the methods from Section 4.2 can be used to construct a confidence set for $\nu^\beta(P_0)$ based on a one-step estimator. This one-step estimator takes the form $\tilde{\nu}_n^\beta := \frac{1}{2} \sum_{j=1}^2 [\Gamma_\beta \circ \nu(\hat{P}_n^j) + P_n^j \phi_n^{j, \beta}]$. By Theorem 3, the main condition for the asymptotic validity of the resulting confidence set is that $\tilde{\nu}_n^\beta$ is an asymptotically linear estimator of $\nu^\beta(P_0)$ with influence function ϕ_0^β . Since $\tilde{\nu}_n^\beta$ is a one-step estimator of $\nu^\beta(P_0)$, rather than $\nu(P_0)$, $\tilde{\nu}_n^\beta$ generally differs from the regularized one-step estimator $\bar{\nu}_n^\beta$ of $\nu(P_0)$ — indeed, $\tilde{\nu}_n^\beta - \bar{\nu}_n^\beta = \frac{1}{2} \sum_{j=1}^2 [\Gamma_\beta \circ \nu(\hat{P}_n^j) - \nu(\hat{P}_n^j)]$. Nevertheless, conditions on the same regularized remainder and drift terms studied to establish rate guarantees for $\bar{\nu}_n^{\beta_n}$ can ensure the asymptotic linearity of $\tilde{\nu}_n^\beta$ — see Corollary S1 in the appendix.

Let $\mathcal{C}_n^\beta(\hat{\zeta}_n)$ denote an asymptotically valid $(1 - \alpha)$ -level confidence set for $\nu^\beta(P_0)$ constructed according to the methods in Section 4.2. Any standardization operator Ω_n satisfying the conditions of Theorem 3 may be used when doing this. For example, if Ω_n is taken to be the identity, then, for a cutoff $\hat{\zeta}_n$ selected via the bootstrap, a spherical confidence set for $\nu^\beta(P_0)$ would take the form

$$\mathcal{C}_n^\beta(\hat{\zeta}_n) := \left\{ h \in \mathcal{H} : \|\tilde{\nu}_n^\beta - h\|_{\mathcal{H}}^2 \leq \hat{\zeta}_n/n \right\}. \quad (29)$$

Since the methods in Section 4.2 require the parameter of interest to be fixed and not depend on sample size, when constructing $\mathcal{C}_n^\beta(\hat{\zeta}_n)$ we require the choice of β to remain fixed as $n \rightarrow \infty$. Handling cases where β changes with n or is selected data-adaptively is an interesting area for future work. To transform the confidence set for $\nu^\beta(P_0)$ into one for $\nu(P_0)$, we take the preimage $\Gamma_\beta^{-1}[\mathcal{C}_n^\beta(\hat{\zeta}_n)] := \{h \in \mathcal{H} : \Gamma_\beta(h) \in \mathcal{C}_n^\beta(\hat{\zeta}_n)\}$. Since $\nu^\beta(P_0) := \Gamma_\beta \circ \nu(P_0)$, this preimage is an asymptotically valid $(1 - \alpha)$ -level confidence set for $\nu(P_0)$ provided $\mathcal{C}_n^\beta(\hat{\zeta}_n)$ is an asymptotically valid $(1 - \alpha)$ -level confidence set for $\nu^\beta(P_0)$. The transformation Γ_β has a left inverse if all entries of β are nonzero, with $\Gamma_\beta^{-1}(h) = \sum_{k=1}^\infty \beta_k^{-1} \langle h, h_k \rangle_{\mathcal{H}} h_k$ for h in the image of Γ_β . Figure 2 illustrates how the map Γ_β^{-1} stretches spherical and Wald-type confidence sets for the regularized parameter $\nu^\beta(P_0)$ into confidence sets for $\nu(P_0)$.

To simplify the discussion, hereafter we focus on the special case where $\mathcal{C}_n^\beta(\hat{\zeta}_n)$ takes the spherical form in (29). In this case, $\Gamma_\beta^{-1}[\mathcal{C}_n^\beta(\hat{\zeta}_n)]$ takes the elliptical form

$$\left\{ h \in \mathcal{H} : \sum_{k=1}^\infty \beta_k^2 \left[\frac{1}{2} \sum_{j=1}^2 \left\{ \langle \nu(\hat{P}_n^j), h_k \rangle_{\mathcal{H}} + P_n^j \dot{\nu}_n^{j,*}(h_k) \right\} - \langle h, h_k \rangle_{\mathcal{H}} \right]^2 \leq \hat{\zeta}_n/n \right\}. \quad (30)$$

Because β_k must tend to zero as $k \rightarrow \infty$ in order for β to belong to ℓ^2 , the $\|\cdot\|_{\mathcal{H}}$ -diameter of this confidence set, namely $\sup_{h, h' \in \Gamma_\beta^{-1}[\mathcal{C}_n^\beta(\hat{\zeta}_n)]} \|h - h'\|_{\mathcal{H}}$, will not converge to zero with sample size. In contrast, the $\|\cdot\|_\beta$ -diameter of this confidence set will generally shrink to zero at an $n^{-1/2}$ -rate, where $\|h\|_\beta := [\sum_{k=1}^\infty \beta_k^2 \langle h, h_k \rangle_{\mathcal{H}}^2]^{1/2}$. Here we note that $\|\cdot\|_\beta$ is a norm on \mathcal{H} if all of the entries of β are nonzero and is otherwise a seminorm.

The confidence set in (30) also satisfies another desirable property, which can be most easily described by studying a corresponding hypothesis test. For fixed $h_0 \in \mathcal{H}$, this test rejects the null hypothesis that $\nu(P_0) = h_0$ in favor of the complementary alternative precisely when $h_0 \notin \Gamma_\beta^{-1}[\mathcal{C}_n^\beta(\hat{\zeta}_n)]$. This test asymptotically controls the type I error at level α when $\Gamma_\beta^{-1}[\mathcal{C}_n^\beta(\hat{\zeta}_n)]$ has asymptotically valid coverage, and, by the triangle inequality, is consistent against fixed alternatives when $\tilde{\nu}_n^\beta$ is a consistent estimator of $\nu^\beta(P_0)$, $\text{plim}_{n \rightarrow \infty} \hat{\zeta}_n < \infty$, and all entries of β are nonzero. We now show that this test also achieves

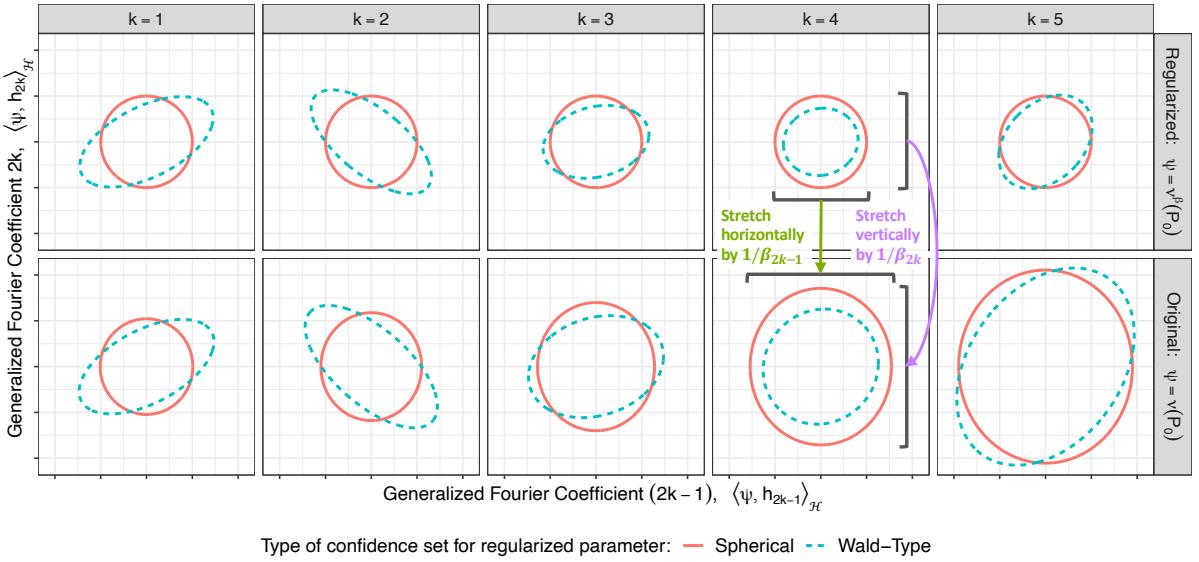


Figure 2: Projections of the boundaries of confidence sets for the regularized parameter $\nu^\beta(P_0)$ (top) and original parameter $\nu(P_0)$ (bottom), plotted via pairs of generalized Fourier coefficients with respect to the basis $(h_k)_{k=1}^\infty$. The transformation applied to confidence sets for $\nu^\beta(P_0)$ to obtain those for $\nu(P_0)$ stretch the axes by the reciprocals of entries of the regularization parameter β . Since $\beta_k \rightarrow 0$ as $k \rightarrow \infty$, this stretch factor becomes arbitrarily large as $k \rightarrow \infty$.

nontrivial power against a class of $n^{-1/2}$ -rate local alternatives. In what follows we let \mathbb{H}^β denote a tight \mathcal{H} -valued Gaussian random variable that is such that, for each $h \in \mathcal{H}$, the marginal distribution $\langle \mathbb{H}^\beta, h \rangle_{\mathcal{H}}$ follows a $N(0, E_0[\langle \phi_0^\beta(Z), h \rangle_{\mathcal{H}}^2])$ distribution, where ϕ_0^β is as defined in Lemma 7. Unlike in the rest of the paper, the following theorem requires all entries of β to be nonzero, since its proof will rely on $\|\cdot\|_\beta$ being a norm.

Theorem 6 (Local power of regularized hypothesis test). *Fix $\beta \in \ell^2 \cap (0, 1]^\mathbb{N}$ and $h_0 \in \mathcal{H}$. Suppose ν is pathwise differentiable at P_0 , $\nu(P_0) = h_0$, $\tilde{\nu}_n^\beta$ is an asymptotically linear estimator of $\nu^\beta(P_0)$ with influence function ϕ_0^β , and $\hat{\zeta}_n$ is a consistent estimator of the $(1 - \alpha)$ -quantile $\zeta_{1-\alpha}$ of $\|\mathbb{H}^\beta\|_{\mathcal{H}}^2$. Fix $\{P_\epsilon : \epsilon\} \in \mathcal{P}(P_0, \mathcal{P}, s)$ such that $\|\dot{\nu}_0(s)\|_{\mathcal{H}} > 0$. If $\Gamma_\beta^{-1}[\mathcal{C}_n^\beta(\hat{\zeta}_n)]$ is as defined in (30), then*

$$P_{\epsilon=n^{-1/2}}^n \left\{ h_0 \notin \Gamma_\beta^{-1}[\mathcal{C}_n^\beta(\hat{\zeta}_n)] \right\} \xrightarrow{n \rightarrow \infty} \Pr \left\{ \|\mathbb{H}^\beta + \dot{\nu}_0^\beta(s)\|_{\mathcal{H}}^2 > \zeta_{1-\alpha} \right\} > \alpha.$$

Also, $h_n := \nu(P_{\epsilon=n^{-1/2}})$ is an $n^{-1/2}$ -rate local alternative in that $\|h_n - h_0\|_{\mathcal{H}} = O(n^{-1/2})$.

The above focuses on local alternatives that are defined via smooth parametric submodels of \mathcal{P} . It is worth noting, however, that by selecting such a submodel, the first-order direction of the local alternative, defined by the value of the local parameter $\dot{\nu}_0(s)$, is fixed as the sample size grows. Since the $\|\cdot\|_{\mathcal{H}}$ -diameter of our confidence set does not decay with sample size, it does not appear that our test will generally have nontrivial asymptotic power against local alternatives whose direction is not fixed and whose $\|\cdot\|_{\mathcal{H}}$ -magnitude decays at an $n^{-1/2}$ rate.

5.3 Tuning parameter selection

We begin by discussing tuning parameter selection for the regularized one-step estimator of $\nu(P_0)$, and then we subsequently discuss confidence set construction. Evaluating the regularized one-step estimator requires selecting three key components: the initial estimator \hat{P}_n^j , orthonormal basis $(h_k)_{k=1}^\infty$, and regularization parameter β_n . Similarly to finite-dimensional problems, the suitability of an initial

estimator \widehat{P}_n^j will depend on the parameter of interest and the form of its corresponding remainder term as defined in (26). In the next section, we will study these remainder terms in our illustrative examples. In what follows we discuss the choice of basis $(h_k)_{k=1}^\infty$ and regularization parameter β_n .

Following the literature on series estimators (Chen, 2007) and motivated by the bound in (28), we suggest choosing the basis $(h_k)_{k=1}^\infty$ so that the span of finitely many initial basis elements yields an accurate approximation of $\nu(P)$, $P \in \mathcal{P}$, provided these Hilbert random elements are smooth enough. Here, smoothness is characterized by the rate of decay of the generalized Fourier coefficients $\langle \nu(P), h_k \rangle_{\mathcal{H}}$ as $k \rightarrow \infty$. If \mathcal{H} is an $L^2(\lambda)$ space with λ the Lebesgue measure on the real line or a bounded subset thereof, then common choices of bases include Legendre polynomials, Laguerre functions, Hermite functions, trigonometric polynomials, and wavelets, among others. If \mathcal{H} is instead an $L^2(Q)$ space with Q an absolutely continuous probability measure on \mathbb{R}^d , then an orthonormal basis for \mathcal{H} can be obtained in several ways. One is to multiply an orthonormal basis $(g_k)_{k=1}^\infty$ for $L^2(\mathbb{R}^d)$ by the root-density $q^{1/2}$ of Q ; in particular, $(q^{1/2}g_k)_{k=1}^\infty$ is an orthonormal basis for $L^2(Q)$. The suitability of these bases for characterizing the smoothness of $\nu(P_0)$ can be assessed on a case-by-case basis. If $d = 1$, then another approach involves transforming an orthonormal basis $(g_k)_{k=1}^\infty$ for $L^2([0, 1])$ via the cumulative distribution function F_Q of Q ; in particular, $(g_k \circ F_Q)_{k=1}^\infty$ is an orthonormal basis of $L^2(Q)$. Other orthonormal bases of $L^2(Q)$ -spaces are also readily available for certain choices of Q , such as if Q is a Gaussian measure (Chapter 9 of Da Prato, 2006). Orthonormal bases for some non- L^2 spaces, such as Sobolev Hilbert spaces, are also well studied (Marcellán and Xu, 2015).

We propose using cross-validation to choose the regularization parameter β_n . If there is uncertainty about which orthonormal basis $(h_k)_{k=1}^\infty$ should be used, this could also be selected via cross-validation, though the discussion that follows focuses on selecting β_n . Our proposal is based on the following loss for $\nu(P_0)$, which relies on an estimate P of P_0 :

$$\mathcal{L}_P(z; h) := \frac{1}{2} \|h - \nu(P)\|_{\mathcal{H}}^2 - \dot{\nu}_P^*[h - \nu(P)](z). \quad (31)$$

The algorithm to implement the proposed cross-validation scheme can be found in Appendix G. There, we also explain why \mathcal{L}_P is a reasonable loss function to use for estimating $\nu(P_0)$. The cross-validation algorithm will be easiest to implement when the search for a regularization parameter is reduced to a search over a finite subset B_n of ℓ_*^2 . A particularly simple choice of B_n consists of the $K_n + 1$ elements of ℓ^2 that take the value 1 in their first $k \in \{0, 1, \dots, K_n\}$ entries and zero in all remaining entries. Selecting over a finite set of possible values is also desirable since there are oracle inequalities for cross-validation selectors over finite sets provided the loss function satisfies appropriate regularity conditions (van der Laan and Dudoit, 2003; van der Vaart et al., 2006). Exploring the applicability of these conditions in our setting is an interesting area for future study.

We now turn to tuning parameter selection for confidence set construction. The considerations for selecting the basis $(h_k)_{k=1}^\infty$ are similar to those discussed above for estimation, and so we focus on selecting the regularization parameter β . As our coverage guarantees rely on the regularization parameter β being fixed and not depending on sample size, cross-validation should not be used to select this quantity. Instead, we recommend choosing a fixed, square-summable sequence β . One natural family of choices is given by setting $\beta = (\beta_k)_{k=1}^\infty$ with $\beta_k = 1/[1 + (k/c)^{1/2+d}]$ for $c, d > 0$. The parameters c and d control the stretch and polynomial rate of decay of the function $k \mapsto \beta_k$, respectively. Finally, we note that, to ensure computational feasibility, the infinite sum used to define the confidence set in (30) can be truncated at a large, finite number of terms K_n^* that grows with n , without adversely affecting coverage. This follows from the fact that the set on the right-hand side of (30) can only be made larger by replacing the sum from $k = 1$ to ∞ with one from $k = 1$ to K_n^* .

6 Study of (regularized) one-step estimators in our examples

We now revisit Examples 1a and 1b from Section 3. For each, we evaluate the plausibility of the regularity conditions that guarantee our theoretical results hold. We revisit the other two examples from Section 3 in Appendices B.3.3 and B.4.4. In what follows, C denotes a generic finite constant whose value may differ from display to display.

Example 1a (Counterfactual density function, continued). Since there is no EIF in this example, we study a regularized one-step estimator $\bar{\nu}_n^{\beta_n}$. This estimator is defined based on an orthonormal basis $(h_k)_{k=1}^\infty$ of \mathcal{H} and a regularization parameter $\beta_n \in \ell_*^2$. Guidance on how to choose these quantities is given in Section 5.3.

Theorem 5 relies on the negligibility of regularized remainder and drift terms $\mathcal{R}_n^{j,\beta_n}$ and $\mathcal{D}_n^{j,\beta_n}$ and bias terms $\mathcal{B}_n^{j,\beta_n}$. In Appendix B.1.3, we use Lemma 5 and the strong positivity assumption to show there exists a constant C that does not depend on β such that, for all $P \in \mathcal{P}$,

$$\|\mathcal{R}_P^\beta\|_{\mathcal{H}} \leq C\|\beta\|_{\ell^2} \|g_P(1 \mid \cdot) - g_0(1 \mid \cdot)\|_{L^2(P_{0,X})} \|p_{Y|A=1,X} - p_{0,Y|A=1,X}\|_{L^2(\tau_0)}, \quad (32)$$

where τ_0 denotes the product measure $\lambda_Y \times P_{0,X}$ and, within the $L^2(\tau_0)$ norm, $p_{Y|A=1,X} - p_{0,Y|A=1,X}$ denotes the function $(y, x) \mapsto p_{Y|A,X}(y \mid 1, x) - p_{0,Y|A,X}(y \mid 1, x)$. The upper bound in (32) depends on three quantities: the ℓ^2 -magnitude of β , the $L^2(P_{0,X})$ -distance between the propensities $g_P(1 \mid \cdot)$ and $g_0(1 \mid \cdot)$, and a root-MISE of the conditional distribution of $Y \mid A = 1, X$ under P relative to that under P_0 , where the mean is taken across values of $X \sim P_0$. Applying the above inequality to study the remainder term $\mathcal{R}_n^{j,\beta_n}$ that Theorem 5 requires to be $O_p[\|\beta_n\|_{\ell^2}/n^{1/2}]$, we see that $\mathcal{R}_n^{j,\beta_n}$ will satisfy this condition provided typical $n^{-1/4}$ -rate conditions are satisfied by the estimators of two nuisance functions, namely the propensity to receive treatment and the conditional density of the outcome given treatment and covariates. Such conditions have been discussed extensively in the literature across a variety of problems (e.g., [van der Laan and Rubin, 2006](#); [Chernozhukov et al., 2018](#)), and tend to hold when the needed nuisance functions are sufficiently smooth or parsimonious relative to the dimension of X and an appropriate estimation strategy is used. For example, suppose that X is continuous and \mathbb{R}^d -valued, $g_0(1 \mid \cdot)$ and $(x, y) \mapsto p_{0,Y|A,X}(y \mid 1, x)$ are Hölder smooth with Hölder exponents b and c , respectively ([Robins et al., 2008](#)). If $g_0(1 \mid \cdot)$ is estimated via a kernel regression and $(x, y) \mapsto p_{0,Y|A,X}(y \mid 1, x)$ is estimated via conditional kernel density estimation, each using kernels of sufficiently high orders, then the above can be used to show that $\mathcal{R}_n^{j,\beta_n} = O_p[\|\beta_n\|_{\ell^2} n^{-\frac{b}{2b+d}} n^{-\frac{c}{2c+d+1}}]$, and so $\mathcal{R}_n^{j,\beta_n}$ achieves the desired $O_p[\|\beta_n\|_{\ell^2}/n^{1/2}]$ rate provided $bc \geq d(d+1)/4$. Alternative estimation strategies that often perform well in practice even when these smoothness assumptions fail, such as those based on random forests ([Ho, 1995](#)) or gradient boosting ([Friedman, 2001](#)), could also be used.

For the regularized drift terms, Lemma 4 shows that $\mathcal{D}_n^{j,\beta_n}$ is $o_p(\|\beta_n\|_{\ell^2}/n^{1/2})$ whenever $\|\phi_n^{j,\beta} - \phi_0^\beta\|_{L^2(P_0; \mathcal{H})} = o_p(\|\beta_n\|_{\ell^2})$. To provide conditions under which this is true, we use that there exists a constant $C > 0$ that does not depend on β such that, for all $P \in \mathcal{P}$, $\|\phi_P^\beta - \phi_0^\beta\|_{L^2(P_0; \mathcal{H})}$ is upper bounded by

$$C\|\beta\|_{\ell^2} \left(\|g_P(1 \mid \cdot) - g_0(1 \mid \cdot)\|_{L^2(P_{0,X})} + \|p_{Y|A=1,X} - p_{0,Y|A=1,X}\|_{L^2(\tau_0)} \right).$$

Hence, $\|\phi_n^{j,\beta_n} - \phi_0^\beta\|_{L^2(P_0; \mathcal{H})} = o_p(\|\beta_n\|_{\ell^2})$ whenever the propensity and conditional density of $Y \mid A = 1, X$ under \hat{P}_n^j are consistent according to the norms above. Consistency is a weaker requirement than the rate conditions imposed to ensure the negligibility of $\mathcal{R}_n^{j,\beta_n}$, so it is reasonable to expect that $\mathcal{D}_n^{j,\beta_n}$ will be negligible when $\mathcal{R}_n^{j,\beta_n}$ is negligible.

From Lemma 6, an upper bound on the rate at which the bias terms $\mathcal{B}_n^{j,\beta_n}$ will decay to zero can be

derived by bounding either $\|\nu(\widehat{P}_n^j) - \nu(P_0)\|_u$ or $\sup_{P \in \mathcal{P}} \|\nu(P)\|_u$ for some $u \geq 0$. The latter of these quantities is no more than $c < \infty$ if the parameter space $\{\nu(P) : P \in \mathcal{P}\}$ is a subset of the Sobolev ellipsoid $\{h \in \mathcal{H} : \|h\|_u \leq c\}$. In this case, when the first K_n entries of β_n are one and all others are zero, Lemma 6 shows that $\|\mathcal{B}_n^{j,\beta_n}\|_{\mathcal{H}} \leq c/(K_n + 1)^u$; if the earlier-discussed regularity conditions hold so that the regularized remainder and drift terms are $O_p(\|\beta_n\|_{\ell^2}/n^{1/2})$, then this yields that, when K_n is of the order $n^{1/(2u+1)}$,

$$\bar{\nu}_n^{\beta_n} - \nu(P_0) = O_p(n^{-u/(2u+1)}). \quad (33)$$

This analysis bears similarity to the study of projection estimators (Theorem 1.9 of [Tsybakov, 2009](#)), but with the added requirement that drift and remainder terms must be considered.

The rate of convergence in (33) was derived based on the looser of the two bounds in (28). While the former bound would give tighter bounds on the bias term when $\|\nu(\widehat{P}_n^j) - \nu(P_0)\|_u$ converges to zero in probability at some rate, it is unclear whether there are initial estimators \widehat{P}_n^j of P_0 that would achieve this. Indeed, since $\|\cdot\|_u$ is a stronger norm than $\|\cdot\|_{\mathcal{H}}$, probabilistic convergence relative to $\|\cdot\|_{\mathcal{H}}$ is insufficient to guarantee convergence relative to $\|\cdot\|_u$. Looking to identify or develop initial estimators of P_0 for which $\|\nu(\widehat{P}_n^j) - \nu(P_0)\|_u \xrightarrow{P} 0$ is an interesting area for future study, since, when such an initial estimator is used, faster rates of convergence for $\bar{\nu}_n^{\beta_n}$ than that given in (33) may be established.

While the discussion above focused on the regularized one-step estimator, similar arguments can be used to analyze the confidence sets introduced in Section 5.2 for fixed $\beta \in \ell_*^2$. Indeed, Corollary S1 in the appendix shows that the key quantities to bound to establish the validity of these confidence sets are $\mathcal{R}_n^{j,\beta}$ and $\mathcal{D}_n^{j,\beta}$ — in particular, both of these quantities should be $o_p(n^{-1/2})$. We have already bounded these quantities above when studying the regularized one-step estimator of $\nu(P_0)$. In particular, these conditions will hold if each of $p_{0,Y|A,X}$ and $g_0(1 \mid \cdot)$ is estimated at a faster-than- $n^{-1/4}$ rate according to the norms in (32).

In the special case where β is a vector whose first K entries are 1 and whose remaining entries are 0, the estimator $\tilde{\nu}_n^{\beta}$ used to construct our confidence sets coincides with the $L^2(\lambda_Y)$ projection estimator studied in Corollary 2 of [Kennedy et al. \(2021\)](#). In our notation, that estimator can be viewed as estimating the parameter $\Gamma_{\beta} \circ \nu(P_0)$, though if K grows with n , as would typically occur under the model selection strategy described by Kennedy et al., then it can be viewed as estimating $\nu(P_0)$ as well. This estimator differs from the regularized one-step estimator $\bar{\nu}_n^{\beta}$ that we have recommended using for estimation of $\nu(P_0)$, with the estimators differing by the $L^2(\lambda_Y)$ projection of $\frac{1}{2} \sum_{j=1}^2 \nu(\widehat{P}_n^j)$ onto the orthogonal complement of the linear span of the first K elements of the chosen basis for $L^2(\lambda_Y)$. It is not immediately clear whether one of these two estimators should be preferred over the other in general, though our upcoming simulation study supports using $\bar{\nu}_n^{\beta}$, especially when n is small. The decision between using these estimators of $\nu(P_0)$ can be summarized as follows: $\bar{\nu}_n^{\beta}$ should be used to estimate $\nu(P_0)$ if $\frac{1}{2} \sum_{j=1}^2 \Gamma_{\beta} \circ \nu(\widehat{P}_n^j)$ attains a lower MISE for estimating $\Gamma_{\beta} \circ \nu(P_0)$ than does the zero function, and $\tilde{\nu}_n^{\beta}$ should be preferred otherwise.

[Kennedy et al. \(2021\)](#) also proposes an approach for making inference about the difference between two counterfactual densities using any of several distance metrics. For the $L^2(\lambda_Y)$ metric, this inference is based upon first-order asymptotics the parameter $\psi(P_0) := \|\nu_1(P_0) - \nu_0(P_0)\|_{L^2(\lambda_Y)}^2$, where ν_1 is equal to the counterfactual density parameter ν defined in (11) and ν_0 takes the same form but with $p_{Y|A,X}(y \mid 1, x)$ replaced by $p_{Y|A,X}(y \mid 0, x)$. There, they note an oft-confronted difficulty ([Luedtke et al., 2019; Williamson et al., 2021](#)) wherein their estimator of $\psi(P_0)$ converges to zero at a faster-than- $n^{-1/2}$ rate under the null hypothesis that $\nu_1(P_0) = \nu_0(P_0)$, leading them to propose a conservative threshold to test this null based on the maximum of the estimated standard error of their estimator and $n^{-1/2}$. Since the pathwise differentiability of ν_1 and ν_0 implies the pathwise differentiability of $\nu_1 - \nu_0$ — with

efficient influence operator equal to the difference of the efficient influence operators of ν_1 and ν_0 — our regularized one-step estimation framework provides an alternative, non-conservative means to test this hypothesis by constructing a confidence set for this parameter for fixed $\beta \in \ell^2 \cap (0, 1]^{\mathbb{N}}$ and checking whether it contains zero.

We now compare our inferential procedure in this example to that of [Kennedy et al. \(2021\)](#). We start by comparing the size of the dual confidence sets. The method from Kennedy et al. can be used to construct a confidence set by inverting tests of whether $\psi_h(P_0) := \|\nu_1(P_0) - \nu_0(P_0) - h\|_{L^2(\lambda_Y)}^2$ is equal to zero across values of $h \in L^2(\lambda_Y)$; the threshold for each h -dependent test is determined using the same conservative threshold methodology as when $h = 0$. The $L^2(\lambda_Y)$ and $\|\cdot\|_\beta$ diameters of this confidence set both decay at rates no faster than $n^{-1/4}$. In contrast, the $\|\cdot\|_\beta$ diameter of our confidence set decays at a quadratically-faster rate of $n^{-1/2}$, while the $L^2(\lambda_Y)$ diameter does not decay at all. Nevertheless, since $\|h\|_\beta = 0$ if and only if $\|h\|_{L^2(\lambda_Y)} = 0$, our confidence set will exclude any particular $h \neq \nu_1(P_0) - \nu_0(P_0)$ with probability tending to one. As a practical matter, our confidence set will exclude functions h that differ smoothly from $\nu_1(P_0) - \nu_0(P_0)$ relative to the basis $(h_k)_{k=1}^\infty$ at smaller sample sizes than will the $n^{-1/4}$ -rate confidence set, and will otherwise require larger sample sizes; here, smoothness is characterized by the decay rate of $\langle \nu_1(P_0) - \nu_0(P_0) - h, h_k \rangle_{\mathcal{H}}$ as $k \rightarrow \infty$. Our dual hypothesis test of whether $\nu_1(P_0) - \nu_0(P_0) = 0$ will also satisfy the local power guarantee from [Theorem 6](#). We investigate the properties of this test and compare it to the test proposed in [Kennedy et al. \(2021\)](#) in our upcoming simulation study.

Example 1b (Bandlimited counterfactual density function, continued). Since there is an EIF in this example, we study a (non-regularized) one-step estimator. [Theorem 2](#) relies on the negligibility of the remainder and drift terms, namely that they are $o_p(n^{-1/2})$. Let $\underline{\mathcal{R}}_P := \underline{\nu}(P) + P_0 \underline{\phi}_P - \underline{\nu}(P_0)$ denote the remainder term for a generic $P \in \mathcal{P}$. In [Appendix B.2.5](#) we show that there exists a $C < \infty$ that does not depend on $P \in \mathcal{P}$ such that

$$\|\underline{\mathcal{R}}_P\|_{\mathcal{H}} \leq C \|g_P(1 \mid \cdot) - g_0(1 \mid \cdot)\|_{L^2(P_0, X)} \|p_{Y|A=1, X} - p_{0, Y|A=1, X}\|_{L^2(\tau_0)}. \quad (34)$$

Hence, for $\underline{\mathcal{R}}_n^j := \underline{\mathcal{R}}_{\hat{P}_n^j}$ to be $o_p(n^{-1/2})$, the products of the rate of convergence of $g_{\hat{P}_n^j}(1 \mid \cdot)$ to $g_0(1 \mid \cdot)$ and $p_{Y|A=1, X}$ to $p_{0, Y|A=1, X}$ according to the norms above must be faster than $n^{-1/2}$. This results in the same $n^{-1/4}$ -type requirement that we discussed below [\(32\)](#) for [Example 1a](#), except, because we only focus on rates of convergence for regularized one-step estimators (rather than weak convergence), there we only required this product to be at least as fast as $n^{-1/2}$ rather than faster, as we require here. Also similarly to [Example 5.3](#), for each $j \in \{1, 2\}$, $\|\underline{\phi}_n^j - \underline{\phi}_0\|_{L^2(P_0; \mathcal{H})}$ can be shown to be $o_p(1)$ provided the propensity and conditional density of $Y \mid A = 1, X$ under \hat{P}_n^j converge to in probability $g_0(1 \mid \cdot)$ and $p_{0, Y|A=1, X}$ according to the norms in [\(34\)](#). Hence, [Lemma 3](#) ensures that the drift term $\underline{\mathcal{D}}_n^j$ is $o_p(n^{-1/2})$ under this condition, and so the conditions of [Theorem 2](#) hold under reasonable conditions. If the operator Ω_0 used to construct a confidence set for $\underline{\nu}(P_0)$ is fixed, then the conditions of [Theorem 4](#) are also satisfied, justifying the use of the bootstrap in confidence set construction. If instead the regularized covariance operator described in [Appendix D](#) is used, then the bootstrap will still yield an asymptotically valid confidence set for $\underline{\nu}(P_0)$ provided the estimator of Ω_0 described in that appendix is used (see [Lemma S12](#) for details).

7 Simulation study

7.1 Overview

We conduct a simulation study to evaluate the finite-sample properties of our one-step estimation framework, both in settings where an EIF exists and in ones where it does not. All of these settings involve drawing inferences about the distributions or densities of counterfactual outcomes (Examples 1 and 3). Our implemented methods are available in the HilbertOneStep R package (Luedtke, 2023).

We consider multiple data-generating processes, each indexed by real-valued probability distributions $Q(0)$ and $Q(1)$. Sampling from a generic such process involves drawing n iid samples from P_0 , where n takes values in $\{250, 500, 1000, 2000, 4000\}$. An observation $Z = (X, A, Y)$ from P_0 is sampled as follows:

$$(Y(0), Y(1)) \sim Q(0) \times Q(1), \quad V | Y(0), Y(1) \sim N(0_5, \text{Id}_5), \\ X = \frac{V}{2} + (V + 1)1\{Y(1) > 0\}; \quad A | Y(0), Y(1), V, X \sim \text{Bern} \left[\frac{1}{20} + \frac{9}{10} \cdot \text{expit}(X_1) \right],$$

and then letting $Y = AY(1) + (1 - A)Y(0)$. For $a \in \{0, 1\}$, $Y(a)$ is the counterfactual outcome if treatment $A = a$ were assigned. Since $Y(a) \perp\!\!\!\perp A | X$ and the positivity assumption is satisfied, the density of $Y(a)$ takes the form in (11) when $a = 1$ and otherwise is the same but with $p_{Y|A,X}(y | 1, x)$ replaced by $p_{Y|A,X}(y | 0, x)$. Unless otherwise specified, estimates of performance are based on 1000 Monte Carlo repetitions.

We estimate all needed nuisance functions using the same approaches as Kennedy et al. (2021). In particular, we estimate the marginal of X with the empirical distribution, the conditional distribution of A given X using the ranger package (Wright and Ziegler, 2017), and the conditional density of Y given (A, X) using ranger and a Gaussian kernel weighted outcome with bandwidth selected by Silverman's rule. When implementing the quadratic forms used to define our confidence sets, we use grids of 500 points on the support of Y — see Appendix F for details.

7.2 Performance of the regularized one-step estimator in Example 1a

We evaluate the performance of the regularized one-step estimator of the counterfactual density of $Y(1)$ from Example 1a when $\mathcal{H} = L^2([0, 1])$. We consider three choices of $Q(1)$, which are displayed in Figure S2 in the appendix. The behavior of $Y(1)$ on the boundaries of its support, namely 0 and 1, differs across the three settings; to emphasize this, we label them ‘zero on both sides’, ‘nonzero on both sides’, and ‘spike on left side’. The cross-validation strategy outlined in Section 5.3 is used to select the regularization parameter β over the elements of ℓ^2 that take the value 1 in their first $K \leq 16$ entries and 0 in all others. To evaluate sensitivity to the choice of basis, we evaluate our estimator based on the cosine basis, with $h_k(y) = 2^{1/2} \cos[\pi(k - 1)y]$, and a rescaled Legendre basis, with $h_k(y)$ proportional to the k -th Legendre function applied to $2y - 1$. We also evaluate the use of cross-validation to select between these bases.

Performance is compared to that of a plugin estimator and also the estimation strategy implemented in the npcausal package (Kennedy et al., 2021), which is a series estimator of the $L^2([0, 1])$ projection of the counterfactual density onto the first K terms of the cosine basis. We use the cross-validation scheme implemented in that package to select a value of $K \leq 16$. Though npcausal does not return the estimated density function, we tweaked its open-source code to extract this information.

Figure 3 displays the estimators' MISEs. In all settings, the regularized one-step estimators are outperformed by the plug-in estimator at small sample sizes, but their relative performances improve as n grows and eventually exceed or trend towards exceeding that of the plugin. Compared to the regularized one-step estimator with the cosine basis, npcausal has MISE that is twice as large at small sample sizes in two of the three scenarios. In one of these scenarios, npcausal's performance improves with n , but is

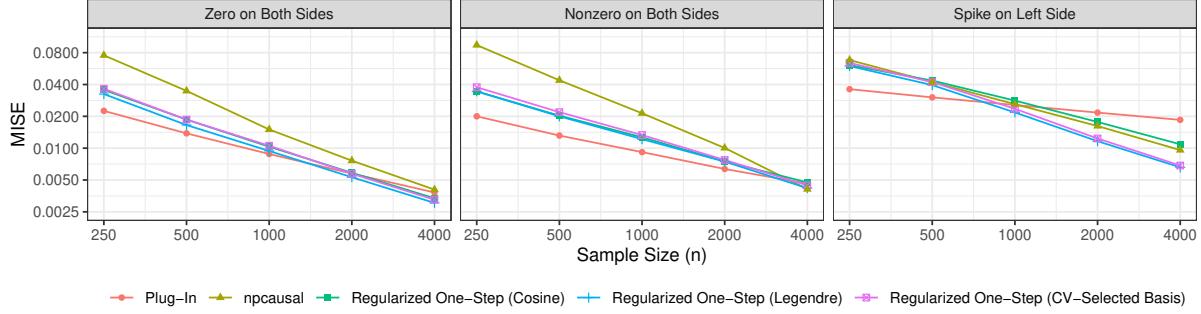


Figure 3: Mean integrated squared error (MISE) versus sample size (n) in Example 1a for five estimators across the three data generating processes considered. Both axes are log-transformed.

still worse than all the other estimators. In the other, its performance dramatically improves between $n = 2000$ to 4000 from the worst of all the estimators to slightly better than the others. In the remaining scenario, npcausal and the regularized one-step estimator with the cosine basis perform similarly. Among the regularized one-step estimators, using the Legendre basis outperforms using the cosine basis in one scenario, while the two perform similarly otherwise. Selecting the basis via cross-validation yields an estimator that is about as good as the one based on the Legendre basis in all scenarios.

7.3 Properties of hypothesis tests from Examples 1a and 3

We evaluate 5% level tests of the null hypothesis that $Q(1) = Q(0)$ against the complementary alternative. The first class of tests uses the results from Example 1a to check if zero is included in an $L^2([0, 1])$ confidence set for the difference $\nu(P_0) = \nu_1(P_0) - \nu_0(P_0)$ of the densities of $Q(1)$ and $Q(0)$. We obtain spherical and Wald-type $L^2([0, 1])$ confidence sets for the regularized parameter $\nu^\beta(P_0)$ using cosine and Legendre bases. Both are then transformed into elliptical confidence sets for $\nu(P_0)$ using the approach from Section 5.2. The Wald-type confidence sets are defined with the correlation-based standardization operator Ω_n from Appendix F with $\lambda = 0.5$. For the regularization parameter β , we let $\beta_k = 1/[1 + (k/c)^2]$ and consider values of $c \in 2.5, 5, 10$. Results for $c = 5$ are reported in the main text, while others appear in the appendix. Additionally, we examine a test based on the Gaussian kernel MMD between $Q(1)$ and $Q(0)$, which depends on a bandwidth choice of 0.5, 1, and 2 times median $\{|Y_i - Y_j| : 1 \leq i < j \leq n\}$ (Garreau et al., 2017). We report results for the middle value in the main text and others in the appendix. We compare performance to the asymptotically conservative test from Kennedy et al. (2021), implemented using npcausal.

We set $Q(1)$ to its value from the ‘nonzero on both sides’ simulation setting and consider different values of $Q(0)$. We explore the null hypothesis with $Q(0) = Q(1)$, and, for $k \in \{1, 2, \dots, 7\}$, the alternative hypothesis with $\nu_0(P_0)(y) = \nu_1(P_0)(y) + \cos(k^2\pi y)/3$, denoted as ‘Alt k^2 ’. By examining these alternatives, we assess the power decay of our tests for $\nu_1(P_0) - \nu_0(P_0) = 0$ as the direction of the alternative corresponds to that of a higher-frequency function in the cosine basis.

Figure 4 presents the type I error and power of the tests. In terms of type I error, the five tests based on our one-step estimation framework achieve or nearly achieve the nominal 5% level. Surprisingly, the npcausal test has a type I error over three times the nominal level at smaller sample sizes, despite having an asymptotic rejection probability of zero. However, as n increases, the type I error converges towards its expected conservative asymptotic behavior. As for power, our tests that regularize using the cosine basis display the anticipated power decay for rejecting Alt k^2 as k increases. Tests using the Legendre basis do not exhibit the same monotonic dependence on k . The Wald-type test with the cosine basis demonstrates noticeably higher power for larger k alternatives than the spherical test, while this trend is less apparent

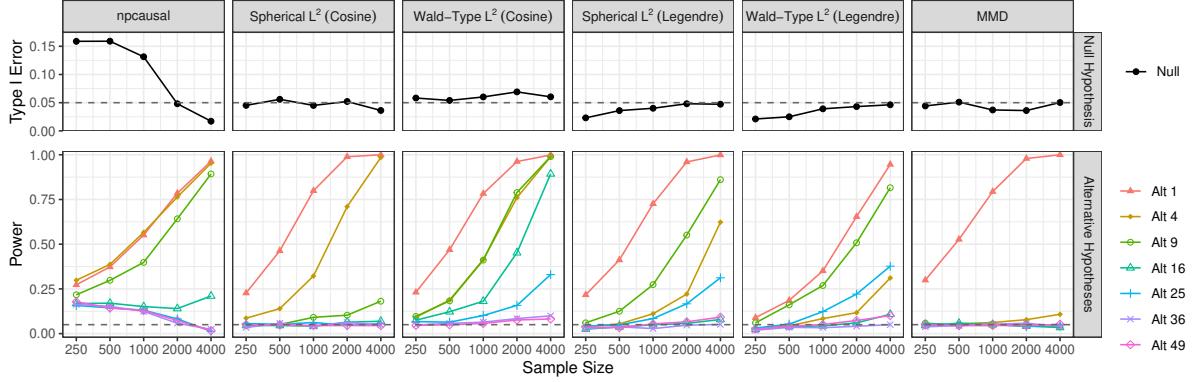


Figure 4: Type I error and power of tests from Examples 1a and 3.

for the Legendre basis. The MMD test has high power for detecting the smoothest alternative, Alt 1, but low or no power against all others. It is important to note that the alternatives we have considered become quite nonsmooth as k increases, making MMD’s poor finite-sample performance for detecting them potentially acceptable. Figure S3 in the appendix illustrates test performance with various tuning parameter choices. Overall, the results align with expectations: for tests based on Example 1a, enhanced power against rougher alternatives (k larger) comes at the cost of reduced power against smoother alternatives when later entries of the regularization parameter β are increased, and vice versa. For MMD, modifying the bandwidth results in the same tradeoff.

Appendix I.1 presents simulation results evaluating our confidence sets for a bandlimited counterfactual density. Nominal coverage is achieved for all sample sizes considered. There, we also highlight the disadvantage of using our L^2 confidence sets when a point evaluation of the counterfactual density, rather than the function itself, is the true quantity of interest.

8 Discussion

The lack of existence of an EIF that we have confronted in parts of this work bears resemblance to the lack of existence of higher-order influence functions for many real-valued parameters (Robins et al., 2008; van der Vaart, 2014; Robins et al., 2017). There, the nonexistence of these objects owes to the lack of a suitable Riesz-representation-type theorem for multilinear forms. In our case, it owes to the fact that the efficient influence operator $\dot{\nu}_P^*(\cdot)(z) : \mathcal{H} \rightarrow \mathbb{R}$ typically fails to be (P -a.s.) bounded and linear, even though $\dot{\nu}_P^* : \mathcal{H} \rightarrow \dot{\mathcal{P}}_P$ is. Though the technical details of the two problems differ, similar solutions work for both: replace the operator that does not satisfy a Riesz-type representation by an approximation that does. In our case, this involved studying the β -regularized efficient influence operator r_P^β . In future work, it would be interesting to investigate the possibility of defining higher-order influence functions for Hilbert-valued parameters. We have shown that, under mild conditions, a first-order EIF exists when the Hilbert space is an RKHS, making this case a natural starting point for exploration.

Another interesting area for future work is to develop a systematic means to select the tuning parameter β needed to define our confidence sets when an EIF does not exist. Although our asymptotic guarantees hold for any fixed choice of $\beta \in \ell_*^2$ and our numerical studies shed light on how different choices of β improve power against different alternatives, it would be desirable to have an automated means of selecting this parameter. One possible approach would involve drawing inspiration from the choice of kernel for two-sample tests based on the MMD (Gretton et al., 2012). Like our approach, these tests rely on selecting a fixed tuning parameter — in their case, a kernel — as n grows. In that context, an appealing heuristic choice of the bandwidth parameter indexing the kernel has been developed, provided

a radial basis function kernel is used (Garreau et al., 2017). It would be of interest to develop a similar heuristic in our setting.

Another area for future study involves extending our results to cases where the Hilbert space depends on the data-generating distribution. As a simple example, in the regression setting from Example 5 in the appendix, we may want to evaluate performance relative to $L^2(\lambda_X)$ with λ_X equal to the marginal distribution $P_{0,X}$ of X under P_0 , rather than some fixed known measure, such as the Lebesgue measure. The definition of pathwise differentiability at P_0 goes through unchanged in that case. However, because the efficient influence operator at an initial estimate \widehat{P}_n^j of P_0 depends on P_0 when $\mathcal{H} = L^2(P_{0,X})$, the regularized one-step estimator we have presented in this work cannot be evaluated. A natural workaround would be to modify the definition of this estimator to use the efficient influence operator of ν at \widehat{P}_n^j relative to a Hilbert space that is indexed by \widehat{P}_n^j , rather than P_0 , and replace $(h_k)_{k=1}^\infty$ by a basis of this space. We leave the study of this estimator to future work.

After establishing that many Hilbert-valued parameters of interest are pathwise differentiable, we focused on developing and studying a particular estimation framework that leverages this property, namely one-step estimation. This framework has the benefit that there is a closed-form expression for the resulting estimators, which simplifies the study of their convergence properties and construction of corresponding confidence sets. While this approach has advantages, it would be worth considering alternative frameworks in future work. As one example, an M-estimator based on the loss that we introduced in (31) could also be considered. Foster and Syrgkanis (2019) offers a general method for determining the convergence rates of these estimators. Additional research is needed to investigate their weak convergence properties and the possibility of using them to construct confidence sets.

Acknowledgements

This work was supported by the National Institutes of Health under award number DP2-LM013340 and the National Science Foundation under award number DMS-2210216. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

References

Agarwal, R., Z. Chen, and S. V. Sarma (2015). Nonparametric estimation of band-limited probability density functions. *arXiv preprint arXiv:1503.06236*.

Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American mathematical society* 68(3), 337–404.

Baiardi, A. and A. A. Naghi (2021). The value added of machine learning to causal inference: Evidence from revisited studies. *arXiv preprint arXiv:2101.00878*.

Berlinet, A. and C. Thomas-Agnan (2011). *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media.

Bickel, P. J., C. A. Klaassen, P. J. Bickel, Y. Ritov, J. Klaassen, J. A. Wellner, and Y. Ritov (1993). *Efficient and adaptive estimation for semiparametric models*, Volume 4. Johns Hopkins University Press Baltimore.

Bogachev, V. (1996). Gaussian measures on linear spaces. *Journal of Mathematical Sciences* 79(2), 933–1034.

Bongioanni, B. and J. L. Torrea (2006). Sobolev spaces associated to the harmonic oscillator. In *Proceedings of the Indian Academy of Sciences-Mathematical Sciences*, Volume 116, pp. 337–360. Springer.

Bongioanni, B. and J. L. Torrea (2008). What is a sobolev space for the laguerre function systems? *Cuadernos de Matemática y Mecánica*.

Cencov, N. N. (1962). Estimation of an unknown distribution density from observations. *Doklady Mathematics*, 1559–1562.

Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. *Handbook of econometrics* 6, 5549–5632.

Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, and W. Newey (2017). Double/debiased/Neyman machine learning of treatment effects. *American Economic Review* 107(5), 261–265.

Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters.

Chernozhukov, V., W. Newey, and R. Singh (2018). De-biased machine learning of global and local parameters using regularized riesz representers. *arXiv preprint arXiv:1802.08667*.

Chernozhukov, V., W. K. Newey, and R. Singh (2021). A simple and general debiased machine learning theorem with finite sample guarantees. *arXiv preprint arXiv:2105.15197*.

Colangelo, K. and Y.-Y. Lee (2020). Double debiased machine learning nonparametric inference with continuous treatments. *arXiv preprint arXiv:2004.03036*.

Da Prato, G. (2006). *An introduction to infinite-dimensional analysis*. Springer Science & Business Media.

Díaz, I. and M. J. van der Laan (2013). Targeted data adaptive estimation of the causal dose–response curve. *Journal of Causal Inference* 1(2), 171–192.

Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *The Annals of Statistics* 7(1), 1–26.

Fawkes, J., R. Hu, R. J. Evans, and D. Sejdinovic (2022). Doubly robust kernel statistics for testing distributional treatment effects even under one sided overlap. *arXiv preprint arXiv:2212.04922*.

Fernique, X. (1970). Intégrabilité des vecteurs gaussiens. *CR Acad. Sci. Paris Serie A* 270, 1698–1699.

Foster, D. J. and V. Syrgkanis (2019). Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*.

Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 1189–1232.

Garreau, D., W. Jitkrittum, and M. Kanagawa (2017). Large sample analysis of the median heuristic. *arXiv preprint arXiv:1707.07269*.

Giné, E. and J. Zinn (1990). Bootstrapping general empirical measures. *The Annals of Probability*, 851–869.

Grenander, U. (1963). *Probabilities on algebraic structures*. Wiley, New York.

Gretton, A., K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola (2012). A kernel two-sample test. *The Journal of Machine Learning Research* 13(1), 723–773.

Hill, J. L. (2011). Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics* 20(1), 217–240.

Hines, O., O. Dukes, K. Diaz-Ordaz, and S. Vansteelandt (2022). Demystifying statistical learning based on efficient influence functions. *The American Statistician*, 1–13.

Ho, T. K. (1995). Random decision forests. In *Proceedings of 3rd international conference on document analysis and recognition*, Volume 1, pp. 278–282. IEEE.

Hudson, A., M. Carone, and A. Shojaie (2021). Inference on function-valued parameters using a restricted score test. *arXiv preprint arXiv:2105.06646*.

Ibragimov, I. and R. Khas’minskii (1983). Estimation of distribution density belonging to a class of entire functions. *Theory of Probability & Its Applications* 27(3), 551–562.

Jung, Y., J. Tian, and E. Bareinboim (2021). Double machine learning density estimation for local treatment effects with instruments. *Advances in Neural Information Processing Systems* 34, 21821–21833.

Kennedy, E. H. (2020). Optimal doubly robust estimation of heterogeneous causal effects. *arXiv preprint arXiv:2004.14497*.

Kennedy, E. H. (2022). Semiparametric doubly robust targeted double machine learning: a review. *arXiv preprint arXiv:2203.06469*.

Kennedy, E. H., S. Balakrishnan, and L. Wasserman (2021). Semiparametric counterfactual density estimation. *arXiv preprint arXiv:2102.12034*.

Kennedy, E. H., S. Balakrishnan, and L. Wasserman (2022). Minimax rates for heterogeneous causal effect estimation. *arXiv preprint arXiv:2203.00837*.

Kennedy, E. H., Z. Ma, M. D. McHugh, and D. S. Small (2017). Non-parametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79(4), 1229–1245.

Klaassen, C. A. (1987). Consistent estimation of the influence function of locally asymptotically linear estimators. *The Annals of Statistics* 15(4), 1548–1562.

Künzel, S. R., J. S. Sekhon, P. J. Bickel, and B. Yu (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences* 116(10), 4156–4165.

Lewandowski, M., M. Ryznar, and T. Źak (1995). Anderson inequality is strict for gaussian and stable measures. *Proceedings of the American Mathematical Society* 123(12), 3875–3880.

Luedtke, A. (2023). Hilbertonestep r package. www.github.com/alexluedtke12/HilbertOneStep.

Luedtke, A., M. Carone, and M. J. van der Laan (2019). An omnibus non-parametric test of equality in distribution for unknown functions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 81(1), 75–99.

Luedtke, A. and J. Wu (2020). Efficient principally stratified treatment effect estimation in crossover studies with absorbent binary endpoints. *Journal de la société française de statistique* 161(1), 176–200.

Luedtke, A. R. and M. J. van der Laan (2016). Super-learning of an optimal dynamic treatment rule. *The international journal of biostatistics* 12(1), 305–332.

Marcellán, F. and Y. Xu (2015). On sobolev orthogonal polynomials. *Expositiones Mathematicae* 33(3), 308–352.

Muandet, K., M. Kanagawa, S. Saengkyongam, and S. Marukatat (2021). Counterfactual mean embeddings. *J. Mach. Learn. Res.* 22, 162–1.

Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. *Handbook of econometrics* 4, 2111–2245.

Nie, X. and S. Wager (2021). Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika* 108(2), 299–319.

Pfanzagl, J. (1982). Lecture notes in statistics. *Contributions to a general asymptotic statistical theory* 13.

Pfanzagl, J. (1990). Estimation in semiparametric models. In *Estimation in Semiparametric Models*, pp. 17–22. Springer.

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(9-12), 1393–1512.

Robins, J., L. Li, R. Mukherjee, E. T. Tchetgen, and A. van der Vaart (2017). Higher order estimating equations for high-dimensional models. *Annals of statistics* 45(5), 1951.

Robins, J., L. Li, E. Tchetgen, A. van der Vaart, et al. (2008). Higher order influence functions and minimax estimation of nonlinear functionals. *Probability and statistics: essays in honor of David A. Freedman* 2, 335–421.

Rudin, W. (1987). *Real and Complex Analysis*. Mathematics series. McGraw-Hill.

Schick, A. (1986). On asymptotically efficient estimation in semiparametric models. *The Annals of Statistics*, 1139–1151.

Schwartz, S. C. (1967). Estimation of probability density by an orthogonal series. *The Annals of Mathematical Statistics*, 1261–1265.

Székely, G. J. and N. K. Bakirov (2003). Extremal probabilities for gaussian quadratic forms. *Probability theory and related fields* 126(2), 184–202.

Takatsu, K. and T. Westling (2022). Debiased inference for a covariate-adjusted regression function. *arXiv preprint arXiv:2210.06448*.

Tikhonov, A. N., A. Goncharsky, V. Stepanov, and A. G. Yagola (1995). *Numerical methods for the solution of ill-posed problems*, Volume 328. Springer Science & Business Media.

Tsiatis, A. A. (2006). Semiparametric theory and missing data.

Tsybakov, A. B. (2009). *Introduction to Nonparametric Estimation*. Springer.

van der Laan, M. J. (2006). Statistical inference for variable importance. *The International Journal of Biostatistics* 2(1).

van der Laan, M. J., A. Bibaut, and A. R. Luedtke (2018). Cv-tmle for nonpathwise differentiable target parameters. In *Targeted Learning in Data Science*, pp. 455–481. Springer.

van der Laan, M. J. and S. Dudoit (2003). Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples.

van der Laan, M. J., M. Laan, and J. M. Robins (2003). *Unified methods for censored longitudinal data and causality*. Springer Science & Business Media.

van der Laan, M. J., S. Rose, et al. (2011). *Targeted learning: causal inference for observational and experimental data*, Volume 10. Springer.

van der Laan, M. J. and D. Rubin (2006). Targeted maximum likelihood learning. *The international journal of biostatistics* 2(1).

van der Vaart, A. (1991). On differentiable functionals. *The Annals of Statistics*, 178–204.

van der Vaart, A. (2014). Higher order tangent spaces and influence functions. *Statistical Science*, 679–686.

van der Vaart, A. and J. A. Wellner (1989). Prohorov and continuous mapping theorems in the hoffmann-jørgensen weak convergence theory, with application to convolution and asymptotic minimax theorems. *Tech. Rep. 157*.

van der Vaart, A. W. (2000). *Asymptotic statistics*, Volume 3. Cambridge university press.

van der Vaart, A. W., S. Dudoit, and M. J. v. d. van der Laan (2006). Oracle inequalities for multi-fold cross validation. *Statistics & Decisions* 24(3), 351–371.

van der Vaart, A. W. and J. Wellner (1996). *Weak convergence and empirical processes: with applications to statistics*. Springer Science & Business Media.

von Mises, R. (1947). On the asymptotic distribution of differentiable statistical functions. *The annals of mathematical statistics* 18(3), 309–348.

Williamson, B. D., P. B. Gilbert, N. R. Simon, and M. Carone (2021). A general framework for inference on algorithm-agnostic variable importance. *Journal of the American Statistical Association*, 1–14.

Wright, M. N. and A. Ziegler (2017). ranger: A fast implementation of random forests for high dimensional data in C++ and R. *Journal of Statistical Software* 77(1), 1–17.

Yao, K. (1967). Applications of reproducing kernel hilbert spaces–bandlimited signal models. *Information and Control* 11(4), 429–444.

Zheng, W. and M. J. Laan (2011). Cross-validated targeted minimum-loss-based estimation. In *Targeted Learning*, pp. 459–474. Springer.

Appendices

Contents

A Additional examples of pathwise differentiable parameters	31
B Derivations for examples	33
B.1 Example 1a: counterfactual density function	33
B.2 Example 1b: bandlimited counterfactual density function	37
B.3 Example 2: counterfactual mean outcome under a continuous treatment	39

B.4	Example 3: counterfactual kernel mean embedding	44
B.5	Example 4: root-density function	47
B.6	Example 5: regression function	47
B.7	Example 6: kernel mean embedding	49
B.8	Example 7: conditional average treatment effect	49
C	Proofs of results from the main text, and supporting lemmas	51
C.1	Proofs for Section 2	51
C.2	Proofs for Section 3	58
C.3	Proofs for Section 4	61
C.4	Proofs for Section 5	67
D	A regularized inverse covariance operator and a consistent estimator thereof	72
E	A conservative estimator of the threshold used to define our confidence sets that does not require the bootstrap	75
F	Numerical considerations for computing the proposed confidence sets	77
G	Cross-validated selection of the regularization parameter	78
H	Inefficient influence operators and influence functions	78
I	Additional simulation results	81
I.1	Simulation results for Example 1b	81
I.2	Supplemental tables and figures for simulation studies	82

A Additional examples of pathwise differentiable parameters

All derivations are given in Appendix B.

Example 4 (Root-density function). Nonparametric density estimation is a well-studied problem. Estimating the (square root of) the density as an L^2 parameter has been done in Cencov (1962). In our setting, we suppose that $Z \sim P$ for $P \in \mathcal{P}$, where there is a σ -finite measure λ that dominates all distributions in a locally nonparametric model \mathcal{P} and we wish to estimate the square root of the density of Z , that is, $\nu(P)(z) := \frac{dP}{d\lambda}(z)^{1/2}$, where λ denotes the Lebesgue measure. The parameter ν takes values in $\mathcal{H} := L^2(\lambda)$. The local parameter takes the form $\dot{\nu}_P(s)(z) = \frac{1}{2}s(z)\nu(P)(z)$, and the efficient influence operator takes the form

$$\dot{\nu}_P^*(h)(z) = \frac{h(z)}{2\nu(P)(z)} - E_P \left[\frac{h(Z)}{2\nu(P)(Z)} \right].$$

Similarly to Example 1a, $\dot{\nu}_P^*$ is not a bounded operator when λ is not discrete, and so an EIF will not exist in these cases.

Example 5 (Regression function). We suppose that $Z := (X, Y) \sim P$ for P in a locally nonparametric model \mathcal{P} . We assume that for all $P, P' \in \mathcal{P}$, P is equivalent to P' in that $P \ll P'$ and $P' \ll P$, and also that

$$\sup_{P \in \mathcal{P}} \text{ess sup}_x E_P[Y^2 | X = x] < \infty, \quad (\text{S1})$$

where ess sup denotes a P_X -essential supremum with P_X denoting the marginal distribution of X under $P \in \mathcal{P}$. Let λ_X be a measure that is dominated by the marginal distribution of X under some (and, by their

equivalence, all) distributions in \mathcal{P} . We wish to estimate the conditional mean $\nu(P)(x) := E_P[Y \mid X = x]$, where $\nu : \mathcal{P} \rightarrow \mathcal{H}$ with $\mathcal{H} := L^2(\lambda_X)$. We establish pathwise differentiability at any $P \in \mathcal{P}$ that is such that $\frac{d\lambda_X}{dP_X}$ is bounded P_X -almost surely. We note that we do not require that $P_X \ll \lambda_X$; in such cases, estimation error quantified via the $L^2(\lambda_X)$ -norm only measures performance on a strict subset of the support of P_X .

The local parameter is $\dot{\nu}_P(s)(x) = \int [y - \nu(P)(x)]s(x, y)P_{Y|X}(dy \mid x)$ and the efficient influence operator is $\dot{\nu}_P^*(h)(x, y) = \frac{d\lambda_X}{dP_X}(x)[y - \nu(P)(x)]h(x)$. Similarly to Example 1a, $\dot{\nu}_P^*$ is not a bounded operator when λ_X is not discrete, and so an EIF will not exist in these cases.

Example 6 (Kernel mean embedding of distributions). Let $\kappa : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ be a bounded, symmetric, positive definite function. Let \mathcal{H} be the unique RKHS associated with the kernel κ (Aronszajn, 1950), and let $K_z := \kappa(z, \cdot)$ be the associated feature map. Assume that \mathcal{P} is a model on \mathcal{Z} such that all $P \in \mathcal{P}$ are equivalent and $P \ll \lambda$ for all $P \in \mathcal{P}$. The target of estimation is the evaluation of the kernel mean embedding $\nu : \mathcal{P} \rightarrow \mathcal{H}$ at P (Gretton et al., 2012), where

$$\nu(P) := \int K_z P(dz).$$

When the model is locally nonparametric, the local parameter takes the form $\dot{\nu}_P(s) = \int K_z s(z) P(dz)$. The efficient influence operator takes the form $\dot{\nu}_P^*(h)(y) = h(y) - Ph$, and the efficient influence function is P -Bochner square integrable and takes the form $\phi_P(z) = K_z - \nu(P)$. Regardless of the initial estimator of P_0 , the one-step estimator is given by $\bar{\nu}_n = \frac{1}{n} \sum_{i=1}^n K_{Z_i}$. In other words, $\bar{\nu}_n$ is the empirical kernel mean embedding as defined in Gretton et al. (2012).

Example 7 (Conditional average treatment effect). There has recently been much interest in various fields regarding the estimation of the conditional average treatment effect function (Hill, 2011; Luedtke and van der Laan, 2016; Künzel et al., 2019). Under conditions, this parameter corresponds to an additive causal effect between the mean outcome that would be observed among individuals with a given covariate value if, possibly contrary to fact, treatment 1 had been administered versus not administered. For the setting, suppose that $Z := (X, A, Y) \sim P$ for $P \in \mathcal{P}$, where X is a vector of covariates, A is a binary treatment, and Y is an outcome. As in Example 4, suppose that \mathcal{P} is a locally nonparametric model consisting of equivalent measures. Define the propensity to receive treatment a as $g_P(a \mid x) := P(A = a \mid X = x)$ and the outcome regression as $\mu_{P,a}(x) := E_P[Y \mid A = a, X = x]$. Suppose that all distributions $P \in \mathcal{P}$ are such that $\max_{a \in \{0,1\}} \text{ess sup}_{x \in \mathcal{X}} E_P(Y^2 \mid A = a, X = x) < \infty$ and $\max_{a \in \{0,1\}} \text{ess sup}_{x \in \mathcal{X}} g_P(a \mid x)^{-1} < \infty$. Let λ_X be a measure that is dominated by the marginal of X under the distributions in \mathcal{P} , and define $\mathcal{H} := L^2(\lambda_X)$. The target of estimation is the conditional average treatment effect $\nu : \mathcal{P} \rightarrow \mathcal{H}$, defined as

$$\nu(P)(x) := \mu_{P,1}(x) - \mu_{P,0}(x).$$

Similarly to Example 4, we establish pathwise differentiability at any $P \in \mathcal{P}$ that is such that $\frac{d\lambda_X}{dP_X}$ is bounded P_X -almost surely. The local parameter takes the form

$$\dot{\nu}_P(s)(x) = \int \frac{(2a - 1)}{g_P(a \mid x)} [y - \mu_{P,a}(x)]s(y, a, x)P(dy, da \mid x). \quad (\text{S2})$$

The efficient influence operator takes the form

$$\dot{\nu}_P^*(h)(y, a, x) = \frac{h(x)}{p_X(x)} \frac{(2a - 1)}{g_P(a \mid x)} [y - \mu_{P,a}(x)]. \quad (\text{S3})$$

Similarly to Example 1a, $\dot{\nu}_P^*$ is not a bounded operator when λ_X is not discrete, and so an EIF will not exist in these cases.

B Derivations for examples

B.1 Example 1a: counterfactual density function

B.1.1 Pathwise differentiability

We now show that ν is pathwise differentiable relative to a locally nonparametric model \mathcal{P} at any $P \in \mathcal{P}$. To do this, we break our argument into two parts. First, we use Lemma 2 to establish that ν is pathwise differentiable relative to a model \mathcal{P}_g that is nonparametric up to the fact that the propensity to receive treatment is known to be equal to a fixed function g . Specifically, we consider the model \mathcal{P}_g that consists of all distributions P' that are such that $g_{P'} = g$ and for which there exists $P \in \mathcal{P}$ such that $P'_{Y|A,X} = P_{Y|A,X}$ and $P'_X = P_X$. Second, we use the fact that ν does not depend on the propensity to receive treatment to extend this pathwise differentiability result to the locally nonparametric model \mathcal{P} .

Throughout this subappendix we suppose that, at each $P \in \mathcal{P}$ and for a fixed $\delta > 0$, \mathcal{P} is large enough to contain submodels of the form $\{P_\epsilon : \epsilon \in [0, \delta)\}$, where $\frac{dP_{\epsilon,X}}{dP_X}(x) = 1 + \epsilon s_X(x)$, $\frac{dP_{\epsilon,A|X}}{dP_{A|X}}(a | x) = 1$, and $\frac{dP_{\epsilon,Y|A,X}}{dP_{Y|A,X}}(y | a, x) = 1 + \epsilon s_{Y|A,X}(y | a, x)$, where s_X and $s_{Y|A,X}$ are arbitrary functions bounded in $(-\delta^{-1}, \delta^{-1})$ that are such that $E_P[s_X(X)] = 0$ and $E_P[s_{Y|A,X}(Y | A, X) | A, X] = 0$ P -almost surely. There is no loss in generality in assuming that \mathcal{P} is this large since pathwise differentiability relative to a larger model also implies pathwise differentiability relative to a smaller model, and, when both the larger and the smaller models are locally nonparametric, the local parameters and efficient influence operators in the two models necessarily agree.

We now use Lemma 2 to prove that ν is pathwise differentiable relative to \mathcal{P}_g , where for now we take g to be any fixed function that is such that $g = g_{P'}$ for at least one $P' \in \mathcal{P}$. Fix two distributions P and \tilde{P} in \mathcal{P}_g . Let λ_X denote a σ -finite measure that dominates the marginals in X of P and \tilde{P} , that is, $P_X \ll \lambda_X$ and $\tilde{P} \ll \lambda_X$, and let q_X and \tilde{q}_X denote the square root of the marginal density of X relative to λ_X under P and \tilde{P} , respectively. For any $P' \in \mathcal{P}$, we also define $q_{P'}(\cdot | x)$ to be the square root of the conditional density of Y given $(A, X) = (1, x)$ under P' . For brevity we let $q := q_P$ and $\tilde{q} := q_{\tilde{P}}$. We have that

$$\begin{aligned}
& \left\| \nu(\tilde{P}) - \nu(P) \right\|_{\mathcal{H}}^2 \\
&= \int \left[\int \{ \tilde{q}^2(y | x) \tilde{q}_X^2(x) - q^2(y | x) q_X^2(x) \} d\lambda_X(x) \right]^2 d\lambda_Y(y) \\
&= \int \left[\int \{ \tilde{q}(y | x) \tilde{q}_X(x) + q(y | x) q_X(x) \right. \\
&\quad \left. \cdot \{ \tilde{q}(y | x) \tilde{q}_X(x) - q(y | x) q_X(x) \} d\lambda_X(x) \right]^2 d\lambda_Y(y) \\
&\leq \int \left[\int \{ \tilde{q}(y | x) \tilde{q}_X(x) + q(y | x) q_X(x) \}^2 d\lambda_X(x) \right] \\
&\quad \cdot \left[\int \{ \tilde{q}(y | x) \tilde{q}_X(x) - q(y | x) q_X(x) \}^2 d\lambda_X(x) \right] d\lambda_Y(y) \\
&\leq 2 \int \left[\nu(\tilde{P})(y) + \nu(P)(y) \right] \left[\int \{ \tilde{q}(y | x) \tilde{q}_X(x) - q(y | x) q_X(x) \}^2 d\lambda_X(x) \right] d\lambda_Y(y),
\end{aligned}$$

where the first inequality holds by Cauchy-Schwarz and the second by the fact that $(b + c)^2 \leq 2(b^2 + c^2)$.

Using (12) and the fact that ν does not depend on the propensity to receive treatment, this shows that, for $C_1 = 4 \sup_{P' \in \mathcal{P}} \text{ess sup}_y \nu(P')(y)$,

$$\begin{aligned} & \left\| \nu(\tilde{P}) - \nu(P) \right\|_{\mathcal{H}}^2 \\ & \leq C_1 \iint \{ \tilde{q}(y | x) \tilde{q}_X(x) - q(y | x) q_X(x) \}^2 d\lambda_X(x) d\lambda_Y(y) \\ & = C_1 \iint \frac{1}{g(1 | x)} \left\{ \tilde{q}(y | x) g^{1/2}(1 | x) \tilde{q}_X(x) - q(y | x) g^{1/2}(1 | x) q_X(x) \right\}^2 d\lambda_X(x) d\lambda_Y(y). \end{aligned}$$

Again using (12) and letting $C_2 = C_1 / \inf_{P' \in \mathcal{P}} \text{ess inf}_{x \in \mathcal{X}} g_{P'}(1 | x)$, we find that

$$\begin{aligned} & \left\| \nu(\tilde{P}) - \nu(P) \right\|_{\mathcal{H}}^2 \\ & \leq C_2 \iint \left\{ \tilde{q}(y | x) g^{1/2}(1 | x) \tilde{q}_X(x) - q(y | x) g^{1/2}(1 | x) q_X(x) \right\}^2 d\lambda_X(x) d\lambda_Y(y). \end{aligned}$$

Finally, noting that the double integral on the right-hand side upper bounds by $H^2(P, \tilde{P})$, we have shown that $\| \nu(\tilde{P}) - \nu(P) \|_{\mathcal{H}} \leq C_2^{1/2} H(P, \tilde{P})$, which establishes (ii) of Lemma 2 when the model is \mathcal{P}_g , where g is an arbitrary value of the propensity to receive treatment for which there exists some $P \in \mathcal{P}$ such that $g = g_P$.

Hereafter we fix $P \in \mathcal{P}$ and suppose that $g = g_P$. We now establish (i) of Lemma 2 at P for the model \mathcal{P}_g with $\eta_P(s)(y)$ equal to the right-hand side of (13) from the main text. To do this, we use the following model: $\{P_\epsilon : \epsilon \in [0, \delta]\}$, where $\frac{dP_{\epsilon, X}}{dP_X}(x) = 1 + \epsilon s_X(x)$, $\frac{dP_{\epsilon, A|X}}{dP_{A|X}}(a | x) = 1$, and $\frac{dP_{\epsilon, Y|A, X}}{dP_{Y|A, X}}(y | a, x) = 1 + \epsilon s_{Y|A, X}(y | a, x)$, where s_X and $s_{Y|A, X}$ are bounded in $[-\delta^{-1}/2, \delta^{-1}/2]$ and $E_P[s_X(X)] = 0$ and $E_P[s_{Y|A, X}(Y | A, X) | A, X] = 0$ P -almost surely. The model $\{P_\epsilon : \epsilon \in [0, \delta]\}$ is a submodel of \mathcal{P} by assumption and, due to the fact that $P_{\epsilon, A|X} = P_{A|X}$ for all ϵ , is therefore also a submodel of \mathcal{P}_g . It can be verified that this submodel has score $s(x, a, y) = s_X(x) + s_{Y|A, X}(y | a, x)$ at $\epsilon = 0$ and that the $L^2(P)$ -closure of the set containing such scores corresponds to the tangent space of \mathcal{P}_g at P . In what follows we will show that $\| \nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s) \|_{\mathcal{H}} = o(\epsilon)$. To this end, observe that

$$\begin{aligned} & \| \nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s) \|_{\mathcal{H}}^2 \\ & = \int \left[\int \left\{ \left(\frac{dP_{\epsilon, Y|A, X}}{dP_{Y|A, X}}(y | 1, x) \frac{dP_{\epsilon, X}}{dP_X}(x) - 1 \right) q^2(y | x) q_X^2(x) \right\} d\lambda_X(x) \right. \\ & \quad \left. - \epsilon \eta_P(s)(y) \right]^2 d\lambda_Y(y) \\ & = \int \left[\int \left\{ (\epsilon s_{Y|A, X}(y | 1, x) \right. \right. \\ & \quad \left. \left. + \epsilon s_X(x) + \epsilon^2 s_{Y|A, X}(y | 1, x) s_X(x)) q^2(y | x) q_X^2(x) \right\} d\lambda_X(x) - \epsilon \eta_P(s)(y) \right]^2 d\lambda_Y(y). \end{aligned}$$

Using that $s_{Y|A, X}(y | a, x) = s(x, a, y) - E_P[s(X, A, Y) | A = a, X = x]$ and $s_X(x) = E_P[s(X, A, Y) | X = x]$ and plugging in the definition of η_P , we see that

$$\begin{aligned} & \| \nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s) \|_{\mathcal{H}}^2 \\ & = \int \left[\int \epsilon^2 s_{Y|A, X}(y | 1, x) s_X(x) q^2(y | x) q_X^2(x) d\lambda_X(x) \right]^2 d\lambda_Y(y) \\ & \leq \epsilon^4 \delta^{-4} \int \left[\int q^2(y | x) q_X^2(x) d\lambda_X(x) \right]^2 d\lambda_Y(y) \end{aligned}$$

$$\begin{aligned}
&= \epsilon^4 \delta^{-4} \int \nu(P)(y)^2 d\lambda_Y(y) \\
&\leq \epsilon^4 \delta^{-4} \left[\int \nu(P)(y) d\lambda_Y(y) \right] \sup_{P \in \mathcal{P}} \text{ess sup}_y \nu(P)(y) \\
&= \epsilon^4 \delta^{-4} \sup_{P \in \mathcal{P}} \text{ess sup}_y \nu(P)(y),
\end{aligned}$$

where the first inequality used that $s_{Y|A,X}$ and s_X both have ranges bounded in $(-\delta^{-1}, \delta^{-1})$. By (12), the right-hand side is $O(\epsilon^4)$, and therefore is $o(\epsilon^2)$ with much to spare. Hence, $\|\nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s)\|_{\mathcal{H}} = o(\epsilon)$.

We now verify that η_P is a bounded operator, which will then show that (i) of Lemma 2 holds at P for the model \mathcal{P}_g . Take any s in the tangent space of \mathcal{P}_g at P . Let $s_{Y|A,X}(y | a, x) := s(x, a, y) - E_P[s(X, A, Y) | A = a, X = x]$ and $s_X(x) := E_P[s(X, A, Y) | X = x]$. It can be verified that, $E_P[s(X, A, Y) | A, X] - E_P[s(X, A, Y) | X] = 0$ P -a.s., and so $s = s_{Y|A,X} + s_X$. We write $p(\cdot | a, x)$ to denote the conditional density under P of Y given $(A, X) = (a, x)$. Observe that

$$\begin{aligned}
&\|\eta_P(s)\|_{L^2(\lambda_Y)}^2 \\
&= \int \left\{ \int [s_{Y|A,X}(y | 1, x) + s_X(x)] p(y | 1, x) P_X(dx) \right\}^2 \lambda_Y(dy) \\
&\leq \int \{ [s_{Y|A,X}(y | 1, x) + s_X(x)] p(y | 1, x) \}^2 \lambda_Y(dy) P_X(dx) \\
&= \int [s_{Y|A,X}(y | 1, x) + s_X(x)]^2 p(y | 1, x) P_{Y|A,X}(dy | 1, x) P_X(dx) \\
&= \int [s_{Y|A,X}(y | 1, x) + s_X(x)]^2 p(y | 1, x) \frac{a}{P(A = 1 | X = x)} P(dz) \\
&= \int [s_{Y|A,X}(y | a, x) + s_X(x)]^2 p(y | a, x) \frac{a}{P(A = 1 | X = x)} P(dz) \\
&\leq \left(\int [s_{Y|A,X}(y | a, x) + s_X(x)]^2 P(dz) \right) \sup_{x,y} \left(p(y | 1, x) \frac{1}{P(A = 1 | X = x)} \right) \\
&\leq \|s\|_{L^2(P)}^2 \sup_{x,y} \left(p(y | 1, x) \frac{1}{P(A = 1 | X = x)} \right).
\end{aligned}$$

By (12), the supremum is finite and so η_P is a bounded linear operator. Hence, by Lemma 2, ν is pathwise differentiable at P relative to \mathcal{P}_g with $\dot{\nu}_P = \eta_P$.

We now show that ν is pathwise differentiable at P relative to \mathcal{P} . To do this, we consider an arbitrary $s \in L_0^2(P)$ and $\{P_\epsilon : \epsilon\} \in \mathcal{P}(P, \mathcal{P}, s)$. Let $s_{Y|A,X}(y | a, x) := s(x, a, y) - E_P[s(X, A, Y) | A = a, X = x]$ and $s_X(x) := E_P[s(X, A, Y) | X = x]$. Also define $\{P'_\epsilon : \epsilon\}$ to be the submodel consisting of distributions P'_ϵ that are such that $P'_{\epsilon,Y|A,X} = P_{\epsilon,Y|A,X}$, $g_{P'_\epsilon} = g_P$, and $P'_{\epsilon,X} = P_{\epsilon,X}$. Lemma S8 can be used to verify that $\{P'_\epsilon : \epsilon\} \in \mathcal{P}(P, \mathcal{P}_g, s_{Y|A,X} + s_X)$. Also, since ν does not depend on the propensity to receive treatment, $\nu(P'_\epsilon) = \nu(P_\epsilon)$ and, by the definition of $\dot{\nu}_P$ in (13), $\dot{\nu}_P(s) = \dot{\nu}_P(s_{Y|A,X} + s_X)$. Hence,

$$\|\nu(P_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s)\|_{L^2(\lambda_Y)} = \|\nu(P'_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s_{Y|A,X} + s_X)\|_{L^2(\lambda_Y)}. \quad (\text{S4})$$

By the pathwise differentiability of ν at P relative to \mathcal{P}_g , the right-hand side is $o(\epsilon)$. Recalling the left-hand side above, this shows that ν is pathwise differentiable at P relative to \mathcal{P} with local parameter $\dot{\nu}_P$.

B.1.2 Efficient influence operator

Take any $s \in L_0^2(P)$. Let $s_{Y|A,X}(y | a, x) := s(x, a, y) - E_P[s(X, A, Y) | A = a, X = x]$ and $s_X(x) := E_P[s(X, A, Y) | X = x]$. To compute the adjoint of $\dot{\nu}_P$, note that, for $h \in L^2(\lambda_Y)$,

$$\langle \dot{\nu}_P(s), h \rangle_{L^2(\lambda_Y)} = \int [s_{Y|A,X}(y | 1, x) + s_X(x)] h(y) P_{Y|A,X}(dy | 1, x) P_X(dx).$$

Note that

$$\begin{aligned} & \int s_{Y|A,X}(y | 1, x) h(y) P_{Y|A,X}(dy | 1, x) P_X(dx) \\ &= \int \frac{1\{a=1\}}{g_P(a | x)} \{h(y) - E_P[h(Y) | A = a, X = x]\} s(z) P(dz), \\ & \int s_X(x) h(y) P_{Y|A,X}(dy | 1, x) P_X(dx) \\ &= \int \{E_P[h(Y) | A = 1, X = x] - E_P E_P[h(Y) | A = 1, X]\} s(z) P(dz). \end{aligned}$$

Thus, (14) holds.

B.1.3 Bounding the regularized remainder term

For $k \in \mathbb{N}$, let $m_{P,k} : x \mapsto E_P[h_k(Y) | A = 1, X = x]$. We now use Lemma 5 to study $\mathcal{R}_n^{j,\beta_n}$. Towards this, note that, for any $P \in \mathcal{P}$ and $k \in \mathbb{N}$,

$$\langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} + P_0 \dot{\nu}_P^*(h_k) = E_0 \left[\left(1 - \frac{g_0(1 | X)}{g_P(1 | X)} \right) (m_{P,k}(X) - m_{0,k}(X)) \right].$$

Hence, by Lemma 5 and the Cauchy-Schwarz inequality,

$$\|\mathcal{R}_P^\beta\|_{\mathcal{H}}^2 \leq \left\| 1 - \frac{g_0(1 | \cdot)}{g_P(1 | \cdot)} \right\|_{L^2(P_0, x)}^2 \sum_{k=1}^{\infty} \beta_k^2 \|m_{P,k} - m_{0,k}\|_{L^2(P_0, x)}^2.$$

Combining this with (12) and letting $C := 1 / \inf_{P \in \mathcal{P}} \text{ess inf}_{x \in \mathcal{X}} g_P(1 | x) < \infty$, we then obtain

$$\|\mathcal{R}_P^\beta\|_{\mathcal{H}}^2 \leq C \|g_P(1 | \cdot) - g_0(1 | \cdot)\|_{L^2(P_0, x)}^2 \sum_{k=1}^{\infty} \beta_k^2 \|m_{P,k} - m_{0,k}\|_{L^2(P_0, x)}^2.$$

Further observe that, for any $k \in \mathbb{N}$, Cauchy-Schwarz and the fact that h_k has unit length in $\mathcal{H} = L^2(\lambda_Y)$ yield that

$$\begin{aligned} & \|m_{P,k} - m_{0,k}\|_{L^2(P_0, x)}^2 \\ &= \int \left(\int h_k(y) [p_{Y|A,X}(y | 1, x) - p_{0,Y|A,X}(y | 1, x)] \lambda_Y(dy) \right)^2 P_{0,X}(dx) \\ &\leq \int \left(\int h_k(y)^2 \lambda_Y(dy) \right) \left(\int [p_{Y|A,X}(y | 1, x) - p_{0,Y|A,X}(y | 1, x)]^2 \lambda_Y(dy) \right) P_{0,X}(dx) \\ &= \iint [p_{Y|A,X}(y | 1, x) - p_{0,Y|A,X}(y | 1, x)]^2 \lambda_Y(dy) P_{0,X}(dx). \end{aligned}$$

Combining the preceding two displays gives (32).

B.2 Example 1b: bandlimited counterfactual density function

B.2.1 Preliminaries

Many of our arguments rely on the following result, which shows the sense in which \mathcal{B} can be viewed as an $L^2(\lambda_Y)$ projection onto $\underline{\mathcal{H}}$. In this lemma and in this lemma only, we are careful to distinguish between elements of $L^2(\lambda_Y)$, which are equivalence classes of functions of the form $[f]$, and elements of $\underline{\mathcal{H}}$, which are functions. After this lemma, we return to following the usual conventions that (i) a generic λ_Y -square integrable function $h : \mathbb{R} \rightarrow \mathbb{R}$ can be treated as an element of $L^2(\lambda_Y)$ by simply replacing h by $[h]$, and (ii) for a generic equivalence class $[f]$ of $L^2(\lambda_Y)$, the function $y \mapsto f(y)$ corresponds to any function belonging to the equivalence class $[f]$. In all contexts where the convention (ii) is used, the particular element of the equivalence class that is selected will not matter.

Lemma S1. *For a λ_Y -square integrable function $h : \mathbb{R} \rightarrow \mathbb{R}$, let $[h]$ denote the equivalence class of functions that are λ_Y -a.e. equal to h . Denoting a generic element of $L^2(\lambda_Y)$ by $[h]$, the operator $[h] \mapsto [\mathcal{B}(h)]$ is an orthogonal projection in $L^2(\lambda_Y)$ onto the closed subspace $\underline{\mathcal{H}}^\dagger := \{[h] : h \in \underline{\mathcal{H}}\}$.*

Proof. Let $\mathcal{B}^\dagger : [h] \mapsto [\mathcal{B}(h)]$. For any function $v : \mathbb{R} \rightarrow \mathbb{C}$, we define $1_{[-b,b]} \cdot v$ as the function $\xi \mapsto 1_{[-b,b]}(\xi) \cdot v(\xi)$.

We first show that \mathcal{B}^\dagger is a linear operator. We show this using that $[f + g] = [f] + [g]$ for any $[f], [g] \in L^2(\lambda_Y)$ and also that the Fourier and inverse Fourier transforms are linear. In particular, for any $[f], [g] \in L^2(\lambda_Y)$ and $c \in \mathbb{R}$, we have that

$$\begin{aligned} \mathcal{B}^\dagger([cf + g]) &= \mathcal{B}^\dagger([cf] + [g]) = [\mathcal{B}(cf + g)] = [\mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F}(cf + g))] \\ &= [\mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F}(cf) + 1_{[-b,b]} \cdot \mathcal{F}(g))] \\ &= [\mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F}(cf)) + \mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F}(g))] \\ &= [c\mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F}(f)) + \mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F}(g))] \\ &= [c\mathcal{B}(f) + \mathcal{B}(g)] = c[\mathcal{B}(f)] + [\mathcal{B}(g)] = c\mathcal{B}^\dagger([f]) + \mathcal{B}^\dagger([g]). \end{aligned}$$

The operator \mathcal{B}^\dagger is idempotent since, for any $[f] \in L^2(\lambda_Y)$,

$$\begin{aligned} \mathcal{B}^\dagger \circ \mathcal{B}^\dagger([f]) &= \mathcal{B}^\dagger([\mathcal{B}(f)]) = [\mathcal{B} \circ \mathcal{B}(f)] \\ &= [\mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F} \circ \mathcal{B}(f))] = [\mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F} \circ \mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F}(f)))] \\ &= [\mathcal{F}^{-1}(1_{[-b,b]} \cdot 1_{[-b,b]} \cdot \mathcal{F}(f))] = [\mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F}(f))] = [\mathcal{B}(f)] = \mathcal{B}^\dagger([f]). \end{aligned}$$

We now show that \mathcal{B}^\dagger is self-adjoint. For any $[f], [g] \in L^2(\lambda_Y)$, the definitions of \mathcal{B} and \mathcal{B}^\dagger show that

$$\begin{aligned} \langle [g], \mathcal{B}^\dagger([f]) \rangle_{L^2(\lambda_Y)} &= \langle [g], [\mathcal{B}(f)] \rangle_{L^2(\lambda_Y)} = \int g(y) \mathcal{B}(f)(y) \lambda_Y(dy) \\ &= \int g(y) \mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F}(f))(y) \lambda_Y(dy) \end{aligned}$$

Applying Plancherel's theorem, the above display continues as follows:

$$\begin{aligned} &= \int \mathcal{F}(g)(y) (1_{[-b,b]} \cdot \mathcal{F}(f))(y) \lambda_Y(dy) \\ &= \int (1_{[-b,b]} \cdot \mathcal{F}(g))(y) \mathcal{F}(f)(y) \lambda_Y(dy) \\ &= \int \mathcal{F}^{-1}(1_{[-b,b]} \cdot \mathcal{F})(g)(y) f(y) \lambda_Y(dy) = \int \mathcal{B}(g)(y) f(y) \lambda_Y(dy) \\ &= \langle [\mathcal{B}(g)], [f] \rangle_{L^2(\lambda_Y)} = \langle \mathcal{B}^\dagger([g]), [f] \rangle_{L^2(\lambda_Y)}. \end{aligned}$$

Hence, \mathcal{B}^\ddagger is an orthogonal projection. Furthermore, the image $\text{Im}(\mathcal{B}^\ddagger)$ of \mathcal{B}^\ddagger can be seen to be equal to $\underline{\mathcal{H}}^\ddagger$. Indeed, (i) $\underline{\mathcal{H}}^\ddagger \subseteq \text{Im}(\mathcal{B}^\ddagger)$ since, for any $[h] \in \underline{\mathcal{H}}^\ddagger$, $[h] = [\mathcal{B}(h)] = \mathcal{B}^\ddagger([h]) \in \text{Im}(\mathcal{B}^\ddagger)$ and (ii) $\text{Im}(\mathcal{B}^\ddagger) \subseteq \underline{\mathcal{H}}^\ddagger$ since, for any $[h] \in \text{Im}(\mathcal{B}^\ddagger)$, the idempotency of \mathcal{B}^\ddagger shows that $\mathcal{F}(h) = \mathcal{F} \circ \mathcal{B}(h) = 1_{[-b,b]} \cdot \mathcal{F}(h)$, and so $h \in \underline{\mathcal{H}}$ and $[h] \in \underline{\mathcal{H}}^\ddagger$. As the image of an orthogonal projection in a Hilbert space is closed, $\underline{\mathcal{H}}^\ddagger$ is a closed subspace of $L^2(\lambda_Y)$. This completes the proof. \square

B.2.2 Pathwise differentiability

Lemma S1 and the pathwise differentiability of ν , established in Example 1a, implies the pathwise differentiability of $\underline{\nu}$ by the following display, which holds for any $\{P_\epsilon : \epsilon \in [0, \delta)\} \in \mathcal{P}(P, \mathcal{P}, s)$:

$$\begin{aligned} \|\underline{\nu}(P_\epsilon) - \underline{\nu}(P_0) - \epsilon \mathcal{B}(\dot{\nu}_P(s))\|_{L^2(\lambda_Y)} &= \|\mathcal{B}(\nu(P_\epsilon) - \nu(P_0) - \epsilon \dot{\nu}_0(s))\|_{L^2(\lambda_Y)} \\ &\leq \|\nu(P_\epsilon) - \nu(P_0) - \epsilon \dot{\nu}_0(s)\|_{L^2(\lambda_Y)} = o(\epsilon). \end{aligned}$$

Hence, $\dot{\underline{\nu}}_P(s) = \mathcal{B} \circ \dot{\nu}_P(s)$. Because \mathcal{B} and $\dot{\nu}_P$ are both linear operators, $\dot{\underline{\nu}}_P : \dot{\mathcal{P}}_P \rightarrow \underline{\mathcal{H}}$ is a linear operator. This operator is also bounded since, by Lemma S1 and the boundedness of $\dot{\nu}_P$, the following holds for any $s \in \dot{\mathcal{P}}_P$:

$$\|\dot{\underline{\nu}}_P(s)\|_{L^2(\lambda_Y)} = \|\mathcal{B} \circ \dot{\nu}_P(s)\|_{L^2(\lambda_Y)} \leq \|\dot{\nu}_P(s)\|_{L^2(\lambda_Y)} \leq \|s\|_{L^2(P)} \|\dot{\nu}_P\|_{\text{op}}.$$

B.2.3 Efficient influence operator

Since \mathcal{B} is self-adjoint, for any $h \in \underline{\mathcal{H}}$,

$$\langle \mathcal{B} \circ \dot{\nu}_P(s), h \rangle_{L^2(\lambda_Y)} = \langle \dot{\nu}_P(s), \mathcal{B}(h) \rangle_{L^2(\lambda_Y)} = \langle s, \dot{\nu}_P^* \circ \mathcal{B}(h) \rangle_{L^2(\lambda_Y)}.$$

Thus, $\dot{\underline{\nu}}_P^*(h) = \dot{\nu}_P^* \circ \mathcal{B}(h)$. Furthermore, as $h \in \underline{\mathcal{H}}$ and \mathcal{B} is a projection onto $\underline{\mathcal{H}}$, $\dot{\underline{\nu}}_P^*(h) = \dot{\nu}_P^*(h)$.

B.2.4 Efficient influence function

To compute the efficient influence function, we let

$$\begin{aligned} \underline{\phi}_P(y, a, x)(\tilde{y}) &:= \dot{\underline{\nu}}_P^*(\underline{K}_{\tilde{y}})(y, a, x) \\ &= \frac{1\{a = 1\}}{g_P(a \mid x)} \left\{ \underline{K}_{\tilde{y}}(y) - E_P [\underline{K}_{\tilde{y}}(Y) \mid A = a, X = x] \right\} \\ &\quad + \left(E_P [\underline{K}_{\tilde{y}}(Y) \mid A = 1, X = x] - \int E_P [\underline{K}_{\tilde{y}}(Y) \mid A = 1, X = \tilde{x}] P_X(d\tilde{x}) \right). \end{aligned}$$

By the symmetry of the kernel, $\underline{K}_{\tilde{y}}(y) = \underline{K}_y(\tilde{y})$. Plugging this into the above and noting that $\underline{\nu}(P) = \int E_P [\underline{K}_Y \mid A = 1, X = \tilde{x}] P_X(d\tilde{x})$, we find that

$$\begin{aligned} \underline{\phi}_P(y, a, x) &= \frac{1\{a = 1\}}{g_P(a \mid x)} \left\{ \underline{K}_y - E_P [\underline{K}_Y \mid A = a, X = x] \right\} + E_P [\underline{K}_Y \mid A = 1, X = x] - \underline{\nu}(P). \end{aligned} \tag{S5}$$

Combining (12) with the square integrability of the sinc function shows that $\underline{\phi}_P$ belongs to $L^2(P; \mathcal{H})$. Hence, Theorem 1 shows that $\underline{\phi}_P$ is the EIF of ν .

B.2.5 Bounding the remainder term

Fix $P \in \mathcal{P}$. Observe that, for any $y \in \mathbb{R}$, the definition of $\underline{\nu}$ and the expression in (S5) yields that

$$\begin{aligned}\underline{\mathcal{R}}_P &:= \underline{\nu}(P) + P_0 \underline{\phi}_P - \underline{\nu}(P_0) \\ &= E_0 \left[\left(1 - \frac{g_0(1 \mid X)}{g_P(1 \mid X)} \right) (E_P[\underline{K}_Y \mid A = 1, X] - E_0[\underline{K}_Y \mid A = 1, X]) \right] \\ &= E_0 \left[\left(1 - \frac{g_0(1 \mid X)}{g_P(1 \mid X)} \right) \int K_y [p_{Y \mid A, X}(y \mid 1, X) - p_{0, Y \mid A, X}(y \mid 1, X)] \lambda_Y(dy) \right].\end{aligned}$$

Hence, letting X_1, X_2 denote independent draws from $P_{0, X}^2$, we find that

$$\begin{aligned}\|\underline{\mathcal{R}}_P\|_{\underline{\mathcal{H}}}^2 &= E_0 \left[\iint K_y(y') \left\{ \prod_{j=1}^2 [p_{Y \mid A, X}(y' \mid 1, X_j) - p_{0, Y \mid A, X}(y' \mid 1, X_j)] \left(1 - \frac{g_0(1 \mid X_j)}{g_P(1 \mid X_j)} \right) \right\} \right. \\ &\quad \left. \lambda_Y(dy) \lambda_Y(dy') \right] \\ &\leq \left(\sup_{y, y' \in \mathbb{R}} |K_y(y')| \right) E_0 \left[\iint \left| \prod_{j=1}^2 [p_{Y \mid A, X}(y' \mid 1, X_j) - p_{0, Y \mid A, X}(y' \mid 1, X_j)] \right. \right. \\ &\quad \left. \left. \cdot \left(1 - \frac{g_0(1 \mid X_j)}{g_P(1 \mid X_j)} \right) \right| \lambda_Y(dy) \lambda_Y(dy') \right].\end{aligned}$$

Using that $\sup_{y, y' \in \mathbb{R}} |K_y(y')| = b/\pi$ and applying Fubini's theorem to the expectation above shows that $\|\underline{\mathcal{R}}_P\|_{\underline{\mathcal{H}}}^2$ is equal to

$$\frac{b}{\pi} E_0 \left[\int \left| [p_{Y \mid A, X}(y' \mid 1, X_1) - p_{0, Y \mid A, X}(y' \mid 1, X_1)] \left(1 - \frac{g_0(1 \mid X_1)}{g_P(1 \mid X_1)} \right) \right| \lambda_Y(dy) \right]^2.$$

Letting $C^2 := b/[\pi \inf_{P' \in \mathcal{P}} \text{ess inf}_{x \in \mathcal{X}} g_{P'}(1 \mid x)^2]$ and applying Cauchy-Schwarz,

$$\|\underline{\mathcal{R}}_P\|_{\underline{\mathcal{H}}}^2 \leq C^2 \|g_P(1 \mid \cdot) - g_0(1 \mid \cdot)\|_{L^2(P_{0, X})}^2 \|p_{Y \mid A=1, X} - p_{0, Y \mid A=1, X}\|_{L^2(\lambda_Y \times P_{0, X})}^2,$$

where $p_{Y \mid A=1, X} - p_{0, Y \mid A=1, X}$ denotes the function $(y, x) \mapsto p_{Y \mid A, X}(y \mid 1, x) - p_{0, Y \mid A, X}(y \mid 1, x)$. We conclude by noting that $C^2 < \infty$ by the strong positivity assumption. Hence, (34) holds.

B.3 Example 2: counterfactual mean outcome under a continuous treatment

B.3.1 Pathwise differentiability

We now show that ν is pathwise differentiable relative to a locally nonparametric model \mathcal{P} at any $P \in \mathcal{P}$. To do this, we follow similar arguments to those used in Appendix B.1.1. In particular, we first use Lemma 2 to establish that ν is pathwise differentiable relative to a model \mathcal{P}_g that is nonparametric up to the fact that the propensity to receive treatment g_P is known to be equal to a fixed function g . Specifically, we consider the model \mathcal{P}_g that consists of all distributions P' that are such that $g_{P'} = g$ and for which there exists $P \in \mathcal{P}$ such that $P'_{Y \mid A, X} = P_{Y \mid A, X}$ and $P'_X = P_X$. Second, we use the fact that ν does not depend on the propensity to receive treatment to extend this pathwise differentiability result to the locally nonparametric model \mathcal{P} .

Let g be such that $g = g_{P'}$ for some fixed $P' \in \mathcal{P}$. We first show that ν is Lipschitz over \mathcal{P}_g . Fix $P, \tilde{P} \in \mathcal{P}_g$. For each $a \in \mathcal{A}$, let P_a and \tilde{P}_a denote the distributions on \mathbb{R} defined so that, for any Borel set

B , $P_a(B) = \int_{\mathcal{X}} \int_B P_{Y|A,X}(dy | a, x) P_X(dx)$ and $\tilde{P}_a(B) = \int_{\mathcal{X}} \int_B \tilde{P}_{Y|A,X}(dy | a, x) \tilde{P}_X(dx)$. We have that

$$\begin{aligned}
& \|\nu(\tilde{P}) - \nu(P)\|_{L^2(\lambda_A)}^2 \\
&= \int \left[\int y \{P_a(dy) - \tilde{P}_a(dy)\} \right]^2 \lambda_A(da) \\
&\leq \int \left[\int y^2 \{P_a^{1/2}(dy) + \tilde{P}_a^{1/2}(dy)\}^2 \right] \left[\int \{P_a^{1/2}(dy) - \tilde{P}_a^{1/2}(dy)\}^2 \right] \lambda_A(da) \\
&\leq \int \left[2 \int y^2 \{P_a(dy) + \tilde{P}_a(dy)\} \right] \left[\int \{P_a^{1/2}(dy) - \tilde{P}_a^{1/2}(dy)\}^2 \right] \lambda_A(da) \\
&\leq \left(2 \sup_{P' \in \mathcal{P}} \text{ess sup}_{a,x} E_{P'}[Y^2 | A = a, X = x] \right) \int \left[\int \{P_a^{1/2}(dy) - \tilde{P}_a^{1/2}(dy)\}^2 \right] \lambda_A(da) \\
&= \left(2 \sup_{P' \in \mathcal{P}} \text{ess sup}_{a,x} E_{P'}[Y^2 | A = a, X = x] \right) \\
&\quad \cdot \int \left[\int \frac{1}{g(a|x)} \{g^{1/2}(a|x) P_a^{1/2}(dy) - g^{1/2}(a|x) \tilde{P}_a^{1/2}(dy)\}^2 \right] \lambda_A(da) \\
&\leq \frac{2 \sup_{P' \in \mathcal{P}} \text{ess sup}_{a,x} E_{P'}[Y^2 | A = a, X = x]}{\inf_{P' \in \mathcal{P}} \text{ess inf}_{(a,x) \in \mathcal{X}} g_{P'}(a|x)} \int [P^{1/2}(dz) - \tilde{P}^{1/2}(dz)]^2 \\
&= \frac{2 \sup_{P' \in \mathcal{P}} \text{ess sup}_{a,x} E_{P'}[Y^2 | A = a, X = x]}{\inf_{P' \in \mathcal{P}} \text{ess inf}_{(a,x) \in \mathcal{X}} g_{P'}(a|x)} H^2(P, \tilde{P}).
\end{aligned}$$

The first inequality above holds by Cauchy-Schwarz, the second by the fact that $(b+c)^2 \leq 2(b^2 + c^2)$, the third by the Hölder's inequality with exponents $(p, q) = (1, \infty)$, and the fourth by the strong positivity assumption. The constant in front of $H^2(P, \tilde{P})$ above is finite by the assumptions regarding the uniform boundedness of the conditional second moment of Y across distributions in \mathcal{P} and the strong positivity assumption. Hence, ν is Lipschitz over \mathcal{P}_g . This establishes (ii) of Lemma 2 when the model is \mathcal{P}_g , where g is an arbitrary value of the propensity to receive treatment for which there exists some $P' \in \mathcal{P}$ such that $g = g_{P'}$.

Hereafter we fix $P \in \mathcal{P}$ and suppose that $g = g_P$. We now establish (i) of Lemma 2 at P for the model \mathcal{P}_g with $\eta_P(s)(a)$ as defined on the right-hand side of (17). To do this, we use the following model: $\{P_\epsilon : \epsilon \in [0, \delta]\}$, where $\frac{dP_{\epsilon,X}}{dP_X}(x) = 1 + \epsilon s_X(x)$, $\frac{dP_{\epsilon,A|X}}{dP_{A|X}}(a|x) = 1$, and $\frac{dP_{\epsilon,Y|A,X}}{dP_{Y|A,X}}(y|a,x) = 1 + \epsilon s_{Y|A,X}(y|a,x)$, where s_X and $s_{Y|A,X}$ are bounded in $[-\delta^{-1}/2, \delta^{-1}/2]$ and $E_P[s_X(X)] = 0$ and $E_P[s_{Y|A,X}(Y|A,X) | A, X] = 0$ P -almost surely. As in Appendix B.1.1, we assume that $\{P_\epsilon : \epsilon \in [0, \delta]\}$ is a submodel of \mathcal{P}_g without loss of generality. This submodel has score $s(x, a, y) = s_X(x) + s_{Y|A,X}(y|a,x)$ at $\epsilon = 0$ and the $L^2(P)$ -closure of the set containing such scores corresponds to the tangent space of \mathcal{P}_g at P . It holds that

$$\begin{aligned}
& \|\nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s)\|_{L^2(\lambda_A)}^2 \\
&= \int [\nu(P_\epsilon)(a) - \nu(P)(a) - \epsilon \eta_P(s)(a)]^2 \lambda_A(da) \\
&= \int \left[\iint y [1 + \epsilon s_{Y|A,X}(y|a,x)] [1 + \epsilon s_X(x)] P_{Y|A,X}(dy|a,x) P_X(dx) \right. \\
&\quad \left. - \iint y P_{Y|A,X}(dy|a,x) P_X(dx) - \epsilon \eta_P(s)(a) \right]^2 \lambda_A(da) \\
&= \epsilon^4 \int \left[\iint y s_{Y|A,X}(y|a,x) s_X(x) P_{Y|A,X}(dy|a,x) P_X(dx) \right]^2 \lambda_A(da)
\end{aligned}$$

$$\begin{aligned}
&\leq \epsilon^4 \delta^{-4} \int \left[\iint |y| P_{Y|A,X}(dy | a, x) P_X(dx) \right]^2 \lambda_A(da) \\
&\leq \epsilon^4 \delta^{-4} \iint \iint y^2 P_{Y|A,X}(dy | a, x) P_X(dx) \lambda_A(da),
\end{aligned}$$

where the first inequality holds by the bounds on the ranges of $s_{Y|A,X}$ and s_X , and the second holds by Jensen's inequality. By the bounds on the conditional second moment of Y under P , the right-hand side is $O(\epsilon^4)$, and so is $o(\epsilon^2)$ with much to spare. Hence, $\|\nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s)\|_{\mathcal{H}} = o(\epsilon)$.

We now verify that η_P is a bounded operator. When combined with the linearity of η_P , this will then show that (i) of Lemma 2 holds at P for the model \mathcal{P}_g . Take any s in the tangent space of \mathcal{P}_g at P . Let $s_{Y|A,X}(y | a, x) := s(x, a, y) - E_P[s(X, A, Y) | A = a, X = x]$ and $s_X(x) := E_P[s(X, A, Y) | X = x]$. It can be verified that, $E_P[s(X, A, Y) | A, X] - E_P[s(X, A, Y) | X] = 0$ P -a.s., and so $s = s_{Y|A,X} + s_X$. Using that $(b + c)^2 \leq 2(b^2 + c^2)$, applying Cauchy-Schwarz and Hölder's inequalities, and leveraging Fubini's theorem, we find that

$$\begin{aligned}
&\|\eta_P(s)\|_{L^2(\lambda_A)}^2 \\
&\leq 2 \int \left[\iint y s_{Y|A,X}(y | a, x) P_{Y|A,X}(dy | a, x) P_X(dx) \right]^2 \lambda_A(da) \\
&\quad + 2 \int \left[\int \mu_P(a, x) s_X(x) P_X(dx) \right]^2 \lambda_A(da) \\
&\leq 2 \int \left[\iint y^2 P_{Y|A,X}(dy | a, x) P_X(dx) \right] \\
&\quad \cdot \left[\iint s_{Y|A,X}^2(y | a, x) P_{Y|A,X}(dy | a, x) P_X(dx) \right] \lambda_A(da) \\
&\quad + 2 \int \left[\int \mu_P^2(a, x) P_X(dx) \right] \left[\int s_X^2(x) P_X(dx) \right] \lambda_A(da) \\
&\leq 2 \left\{ \text{ess sup}_{a,x} E_P[Y^2 | A = a, X = x] \right\} \\
&\quad \cdot \iint \left[\int s_{Y|A,X}^2(y | a, x) P_{Y|A,X}(dy | a, x) + s_X^2(x) \right] P_X(dx) \lambda_A(da) \\
&= 2 \left\{ \text{ess sup}_{a,x} E_P[Y^2 | A = a, X = x] \right\} \\
&\quad \cdot \iint \frac{1}{g_P(a | x)} \left[\int s_{Y|A,X}^2(y | a, x) P_{Y|A,X}(dy | a, x) + s_X^2(x) \right] \\
&\quad \cdot P_X(dx) g_P(a | x) \lambda_A(da) \\
&\leq 2 \frac{\text{ess sup}_{a,x} E_P[Y^2 | A = a, X = x]}{\text{ess inf}_{a,x} g_P(a | x)} \\
&\quad \cdot \iint \left[\int s_{Y|A,X}^2(y | a, x) P_{Y|A,X}(dy | a, x) + s_X^2(x) \right] P_X(dx) g_P(a | x) \lambda_A(da) \\
&= 2 \frac{\text{ess sup}_{a,x} E_P[Y^2 | A = a, X = x]}{\text{ess inf}_{a,x} g_P(a | x)} \int \left[s_{Y|A,X}^2(y | a, x) + s_X^2(x) \right] P(dz) \\
&= 2 \frac{\text{ess sup}_{a,x} E_P[Y^2 | A = a, X = x]}{\text{ess inf}_{a,x} g_P(a | x)} \|s\|_{L^2(P)}^2.
\end{aligned}$$

Above all essential suprema and infima are under the joint distribution of (A, X) implied by P . The fraction above is finite by the strong positivity assumption and the assumed bound on the conditional second moment of Y . Hence, η_P is a bounded operator. By Lemma 2, ν is pathwise differentiable at P relative to \mathcal{P}_g with $\dot{\nu}_P = \eta_P$. In the same way as was done in (S4) for Example 1a, this pathwise

differentiability over \mathcal{P}_g can be extended to show that ν is pathwise differentiable over the locally nonparametric model \mathcal{P} .

B.3.2 Efficient influence operator

Take any $s \in L_0^2(P)$. Let $s_{Y|A,X}(y | a, x) := s(y, a, x) - E_P[s(X, A, Y) | A = a, X = x]$ and $s_X(x) := E_P[s(X, A, Y) | X = x]$. To compute the adjoint of $\dot{\nu}_P$, note that, for any $h \in L^2(\lambda_A)$,

$$\begin{aligned} \langle h, \dot{\nu}_P(s) \rangle_{L^2(\lambda_A)} &= \iiint \{y - \mu_P(a, x)\} s_{Y|A,X}(y, a, x) h(a) P_{Y|A,X}(dy | a, x) P_X(dx) \lambda_A(da) \\ &\quad + \iint [\mu_P(a, x) - \nu(P)(a)] h(a) s_X(x) P_X(dx) \lambda_A(da). \end{aligned}$$

The first term may be rearranged as follows:

$$\begin{aligned} &\iiint \{y - \mu_P(a, x)\} s_{Y|A,X}(y, a, x) h(a) P_{Y|A,X}(dy | a, x) P_X(dx) \lambda_A(da) \\ &= \int \frac{y - \mu_P(a, x)}{g_P(a | x)} h(a) s_{Y|A,X}(y | a, x) P(dx) \\ &= \int \frac{y - \mu_P(a, x)}{g_P(a | x)} h(a) s(z) P(dz). \end{aligned}$$

For the second term:

$$\begin{aligned} &\iint [\mu_P(a, x) - \nu(P)(a)] h(a) s_X(x) P_X(dx) \lambda_A(da) \\ &= \int \left[\int [\mu_P(a, x) - \nu(P)(a)] h(a) \lambda_A(da) \right] s_X(x) P_X(dx) \\ &= \int \left[\int [\mu_P(a', x) - \nu(P)(a')] h(a') \lambda_A(da') \right] s(z) P(dz). \end{aligned}$$

Thus,

$$\dot{\nu}_P^*(h)(y, a, x) = \frac{y - \mu_P(a, x)}{g_P(a | x)} h(a) + \int [\mu_P(a', x) - \nu(P)(a')] h(a') d\lambda_A(a').$$

B.3.3 Study of regularized one-step estimator

Since there is no EIF in this example, we study a regularized one-step estimator $\bar{\nu}_n^{\beta_n}$. This estimator is defined based on an orthonormal basis $(h_k)_{k=1}^\infty$ — guidance for selecting this basis is given in Section 5.3. We study the regularized remainder, regularized drift, and bias terms appearing in Theorem 5 and establish a rate of convergence of $\bar{\nu}_n^{\beta_n}$ for an appropriately chosen sequence of regularization parameters β_n . In what follows, C denotes a generic finite constant whose value may differ from display to display.

We begin by bounding the regularized remainder terms. We use Lemma 5 to derive our bound. To this end, we note that, for any $P \in \mathcal{P}$ and $k \in \mathbb{N}$,

$$\begin{aligned} &\langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} + P_0 \dot{\nu}_P^*(h_k) \\ &= E_0 \left[\frac{Y - \mu_P(A, X)}{g_P(A | X)} h_k(A) + \int [\mu_P(a, X) - \nu(P_0)(a)] h_k(a) \lambda_A(da) \right] \\ &= E_0 \left[\frac{\mu_0(A, X) - \mu_P(A, X)}{g_P(A | X)} h_k(A) + \int [\mu_P(a, X) - \mu_0(a, X)] h_k(a) \lambda_A(da) \right] \\ &= E_0 \left[\int \left[1 - \frac{g_0(a | X)}{g_P(a | X)} \right] [\mu_P(a, X) - \mu_0(a, X)] h_k(a) \lambda_A(da) \right]. \end{aligned}$$

From here, different bounds are possible, depending on the basis $(h_k)_{k=1}^\infty$. If the functions in h_k are

uniformly bounded — that is, $\sup_{a \in \mathcal{A}, k \in \mathbb{N}} |h_k(a)| < \infty$ — then the strong positivity assumption and the Cauchy-Schwarz inequality together show that there is a finite constant C that does not depend on $P \in \mathcal{P}$ or $k \in \mathbb{N}$ such that

$$|\langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} + P_0 \dot{\nu}_P^*(h_k)| \leq C \|g_P - g_0\|_{L^2(\lambda_A \times P_0, X)} \|\mu_P - \mu_0\|_{L^2(\lambda_A \times P_0, X)},$$

where, for $f : (a, x) \mapsto \mathbb{R}$ and $q \geq 1$, $\|f\|_{L^q(\lambda_A \times P_0, X)}^q := \iint f(a, x)^q \lambda_A(da) P_0(dx)$ and $g_P - g_0$ denotes the function $(a, x) \mapsto g_P(a | x) - g_0(a | x)$. Requiring functions in $(h_k)_{k=1}^\infty$ to be uniformly bounded is not such a strong condition, with this condition being satisfied by both the trigonometric and cosine bases for $L^2([0, 1])$. If a basis is used that does not satisfy this assumption, then the following alternative bound can be derived by twice applying Cauchy-Schwarz, using Jensen's inequality, invoking the strong positivity assumption, and leveraging the fact that all elements of $(h_k)_{k=1}^\infty$ have unit length in $L^2([0, 1])$: there exists a C that does not depend on $P \in \mathcal{P}$ or $k \in \mathbb{N}$ such that

$$|\langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} + P_0 \dot{\nu}_P^*(h_k)| \leq C \|g_P - g_0\|_{L^4(\lambda_A \times P_0, X)} \|\mu_P - \mu_0\|_{L^4(\lambda_A \times P_0, X)}.$$

Plugging the above bounds into Lemma 5 shows that, depending on whether or not the functions in $(h_k)_{k=1}^\infty$ are uniformly bounded ($q = 2$) or not ($q = 4$), there exists a $C < \infty$ such that

$$\|\mathcal{R}_P^\beta\|_{\mathcal{H}} \leq C \|\beta\|_{\ell^2} \|g_P - g_0\|_{L^q(\lambda_A \times P_0, X)} \|\mu_P - \mu_0\|_{L^q(\lambda_A \times P_0, X)}.$$

Hence, for each $j \in \{1, 2\}$, $\mathcal{R}_n^{j, \beta_n} := \mathcal{R}_{\hat{P}_n^j}^{\beta_n}$ will be $O_p(\|\beta_n\|_{\ell^2}/n^{1/2})$ provided the product of the rates of convergence in probability of $g_{\hat{P}_n^j}$ and $\mu_{\hat{P}_n^j}$ to g_0 and μ_0 under the $L^q(\lambda_A \times P_0, X)$ norm is at least $n^{-1/2}$.

We now turn to the regularized drift terms. We will bound them via Lemma 4. We begin by noting that, for any $P \in \mathcal{P}$ and $\beta \in \ell_*^2$,

$$\begin{aligned} \|\phi_P^\beta - \phi_0^\beta\|_{L^2(P_0; \mathcal{H})}^2 &= \sum_{k=1}^\infty \beta_k^2 \|\dot{\nu}_P(h_k) - \dot{\nu}_0(h_k)\|_{L^2(P_0)}^2 \\ &\leq \|\beta\|_{\ell^2}^2 \sup_{k \in \mathbb{N}} \|\dot{\nu}_P(h_k) - \dot{\nu}_0(h_k)\|_{L^2(P_0)}^2. \end{aligned}$$

It can further be shown that there exists a constant $C < \infty$ that does not depend on $P \in \mathcal{P}$ or $k \in \mathbb{N}$ such that

$$\|\dot{\nu}_P(h_k) - \dot{\nu}_0(h_k)\|_{L^2(P_0)} \leq C (\|g_P - g_0\|_{L^2(\lambda_A \times P_0, X)} + \|\mu_P - \mu_0\|_{L^2(\lambda_A \times P_0, X)}).$$

Combining the preceding two displays shows that

$$\|\phi_P^\beta - \phi_0^\beta\|_{L^2(P_0; \mathcal{H})} \leq C \|\beta\|_{\ell^2} (\|g_P - g_0\|_{L^2(\lambda_A \times P_0, X)} + \|\mu_P - \mu_0\|_{L^2(\lambda_A \times P_0, X)}).$$

Hence, for each $j \in \{1, 2\}$, $\|\phi_n^{j, \beta_n} - \phi_0^{\beta_n}\|_{L^2(P_0; \mathcal{H})} = o_p(\|\beta_n\|_{\ell^2})$ provided $g_{\hat{P}_n^j}$ and $\mu_{\hat{P}_n^j}$ converge in probability to g_0 and μ_0 under the $L^2(\lambda_A \times P_0, X)$ norm. Since no requirement is made on the rate of convergence and the $L^4(\lambda_A \times P_0, X)$ is stronger than the $L^2(\lambda_A \times P_0, X)$ norm, this condition will typically be weaker than the condition required above to make the regularized remainder term negligible. In any case, under this consistency condition, Lemma 4 shows that $\mathcal{D}_n^{j, \beta_n}$ is $o_p(\|\beta_n\|_{\ell^2}/n^{1/2})$, as desired.

The analysis of the bias term is nearly identical to the one given for Example 1a in the main text. In particular, if the first K_n entries of β_n are one and all the others are zero, then, provided $\sup_{P \in \mathcal{P}} \|\nu(P)\|_u < \infty$, Lemma 6 shows that $\|\mathcal{B}_n^{j, \beta_n}\|_{\mathcal{H}} \leq c/(K_n + 1)^u$. Hence, if K_n is of the order

$n^{1/(2u+1)}$ and the regularized remainder and drift terms are $O_p(\|\beta_n\|_{\ell^2}/n^{1/2})$, then

$$\bar{\nu}_n^{\beta_n} - \nu(P_0) = O_p(n^{-u/(2u+1)}).$$

B.4 Example 3: counterfactual kernel mean embedding

B.4.1 Pathwise differentiability

We now show that ν is pathwise differentiable relative to a locally nonparametric model \mathcal{P} at any $P \in \mathcal{P}$. To do this, we follow similar arguments to those used in Appendix B.1.1. In particular, first we use Lemma 2 to establish that ν is pathwise differentiable relative to the same model \mathcal{P}_g considered in Appendix B.1.1. Second, we use the fact that ν does not depend on the propensity to receive treatment to extend this pathwise differentiability result to the locally nonparametric model \mathcal{P} .

Let g be such that $g = g_{P'}$ for some fixed $P' \in \mathcal{P}$. We first show that ν is Lipschitz over \mathcal{P}_g . Fix $P, \tilde{P} \in \mathcal{P}_g$. For each $a \in \mathcal{A}$, let P_a and \tilde{P}_a denote the distributions on \mathbb{R} defined so that, for any Borel set B , $P_1(B) = \int_{\mathcal{X}} \int_B P_{Y|A,X}(dy \mid 1, x) P_X(dx)$ and $\tilde{P}_1(B) = \int_{\mathcal{X}} \int_B \tilde{P}_{Y|A,X}(dy \mid 1, x) \tilde{P}_X(dx)$. Observe that

$$\begin{aligned} & \|\nu(P) - \nu(\tilde{P})\|_{\mathcal{H}}^2 \\ &= \iint \kappa(y_1, y_2) \prod_{i=1}^2 (P_1 - \tilde{P}_1)(dy_i) \\ &= \iint \kappa(y_1, y_2) \prod_{i=1}^2 \frac{a_i}{g(a_i \mid x_i)} (P - \tilde{P})(dz_i) \\ &= \iint \kappa(y_1, y_2) \prod_{i=1}^2 \frac{a_i}{g(a_i \mid x_i)} \left[\sqrt{dP(z_i)} + \sqrt{d\tilde{P}(z_i)} \right] \left[\sqrt{dP(z_i)} - \sqrt{d\tilde{P}(z_i)} \right] \\ &\leq \left(\iint \kappa^2(y_1, y_2) \prod_{i=1}^2 \frac{a_i}{g^2(a_i \mid x_i)} \left[\sqrt{dP(z_i)} + \sqrt{d\tilde{P}(z_i)} \right]^2 \right)^{1/2} \\ &\quad \cdot \left(\iint \prod_{i=1}^2 \left[\sqrt{dP(z_i)} - \sqrt{d\tilde{P}(z_i)} \right]^2 \right)^{1/2}, \end{aligned}$$

where the inequality holds by Cauchy-Schwarz. The latter of the two terms in the product on the right-hand side is equal to $H^2(P, \tilde{P})$. Using the inequality $(b+c)^2 \leq 2(b^2 + c^2)$ and then applying Hölder's inequality with exponents $(p, q) = (1, \infty)$, the square of the former term in this product bounds as follows:

$$\begin{aligned} & \iint \kappa^2(y_1, y_2) \prod_{i=1}^2 \frac{a_i}{g^2(a_i \mid x_i)} \left[\sqrt{dP(z_i)} + \sqrt{d\tilde{P}(z_i)} \right]^2 \\ &\leq 2 \iint \kappa^2(y_1, y_2) \prod_{i=1}^2 \frac{a_i}{g^2(a_i \mid x_i)} (P + \tilde{P})(dz_i) \\ &\leq \frac{8 \sup_{y_1, y_2 \in \mathcal{Y}} \kappa^2(y_1, y_2)}{\inf_{P' \in \mathcal{P}} \text{ess inf}_x g_{P'}^2(1 \mid x)}. \end{aligned}$$

The right-hand side above is finite by the strong positivity assumption and the fact that κ is bounded. Hence, ν is Lipschitz over \mathcal{P}_g . Combining the preceding two displays establishes (ii) of Lemma 2 when the model is \mathcal{P}_g , where g is an arbitrary value of the propensity to receive treatment for which there exists some $P' \in \mathcal{P}$ such that $g = g_{P'}$.

Hereafter we fix $P \in \mathcal{P}$ and suppose that $g = g_P$. We now establish (i) of Lemma 2 at P for the model \mathcal{P}_g with $\eta_P(s)$ as defined on the right-hand side of (19). To do this, we use the following

model: $\{P_\epsilon : \epsilon \in [0, \delta]\}$, where $\frac{dP_{\epsilon,X}}{dP_X}(x) = 1 + \epsilon s_X(x)$, $\frac{dP_{\epsilon,A|X}}{dP_{A|X}}(a | x) = 1$, and $\frac{dP_{\epsilon,Y|A,X}}{dP_{Y|A,X}}(y | a, x) = 1 + \epsilon s_{Y|A,X}(y | a, x)$, where s_X and $s_{Y|A,X}$ are bounded in $[-\delta^{-1}/2, \delta^{-1}/2]$ and $E_P[s_X(X)] = 0$ and $E_P[s_{Y|A,X}(Y | A, X) | A, X] = 0$ P -almost surely. As in Appendix B.1.1, we assume that $\{P_\epsilon : \epsilon \in [0, \delta]\}$ is a submodel of \mathcal{P}_g without loss of generality. This submodel has score $s(x, a, y) = s_X(x) + s_{Y|A,X}(y | a, x)$ at $\epsilon = 0$ and the $L^2(P)$ -closure of the set containing such scores corresponds to the tangent space of \mathcal{P}_g at P . It holds that

$$\begin{aligned} & \|\nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s)\|_{\mathcal{H}}^2 \\ &= \epsilon^4 \left\| \iint K_y s_{Y|A,X}(y | a, x) s_X(x) P_{Y|A,X}(dy | 1, x) P_X(dx) \right\|_{\mathcal{H}}^2 \\ &= \epsilon^4 \left\| \int \frac{a}{g_P(a | x)} K_y s_{Y|A,X}(y | a, x) s_X(x) P(dz) \right\|_{\mathcal{H}}^2. \end{aligned}$$

The right-hand side is certainly $o(\epsilon^2)$ if the squared \mathcal{H} -norm on that side is finite. To see that this is the case, note first that, by the strong positivity assumption and the fact that κ , $s_{Y|A,X}$, and s_X are all bounded functions, $(x, a, y) \mapsto \frac{a}{g_P(a | x)} K_y s_{Y|A,X}(y | a, x) s_X(x)$ belongs to $L^2(P; \mathcal{H})$. Hence, that term satisfies the following:

$$\begin{aligned} & \left\| \int \frac{a}{g_P(a | x)} K_y s_{Y|A,X}(y | a, x) s_X(x) P(dz) \right\|_{\mathcal{H}}^2 \\ &= \iint \frac{a}{g_P(a | x)} \frac{a'}{g_P(a' | x')} \kappa(y, y') s_{Y|A,X}(y | a, x) s_X(x) P(dz) P(dz') < \infty, \end{aligned}$$

where $z' = (x', a', y')$. This establishes that $\|\nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s)\|_{\mathcal{H}} = o(\epsilon)$.

We now verify that η_P is a bounded operator. When combined with the linearity of η_P , this will then show that (i) of Lemma 2 holds at P for the model \mathcal{P}_g . Take any s in the tangent space of \mathcal{P}_g at P . Let $s_{Y|A,X}(y | a, x) := s(x, a, y) - E_P[s(X, A, Y) | A = a, X = x]$ and $s_X(x) := E_P[s(X, A, Y) | X = x]$. It can be verified that, $E_P[s(X, A, Y) | A, X] - E_P[s(X, A, Y) | X] = 0$ P -a.s., and so $s = s_{Y|A,X} + s_X$. Since s is P -square integrable, $s_{Y|A,X}$ and s_X are as well. By rewriting the right-hand side of (19), we see that η_P satisfies:

$$\eta_P(s) = \int \frac{a}{g_P(a | x)} K_y [s_{Y|A,X}(y | a, x) + s_X(x)] P(dz). \quad (\text{S6})$$

By the strong positivity assumption, the fact that κ is a bounded function, and the fact that $s_{Y|A,X}$ and s_X belong to $L^2(P)$, $(x, a, y) \mapsto \frac{a}{g_P(a | x)} K_y [s_{Y|A,X}(y | a, x) + s_X(x)]$ belongs to $L^2(P; \mathcal{H})$. Hence,

$$\begin{aligned} \|\eta_P(s)\|_{\mathcal{H}}^2 &= \iint \frac{a}{g_P(a | x)} \frac{a'}{g_P(a' | x')} \kappa(y, y') [s_{Y|A,X}(y | a, x) + s_X(x)] \\ &\quad \cdot [s_{Y|A,X}(y' | a', x') + s_X(x')] P^2(dz, dz') \\ &\leq \iint \frac{a}{g_P(a | x)} \frac{a'}{g_P(a' | x')} \sqrt{\kappa(y, y') \kappa(y', y')} |s_{Y|A,X}(y | a, x) + s_X(x)| \\ &\quad \cdot |s_{Y|A,X}(y' | a', x') + s_X(x')| P^2(dz, dz') \\ &= \left[\int \frac{a}{g_P(a | x)} \sqrt{\kappa(y, y)} |s_{Y|A,X}(y | a, x) + s_X(x)| P(dz) \right]^2 \\ &\leq \left[\int \frac{a}{g_P^2(a | x)} |\kappa(y, y)| P(dz) \right] \left[\int [s_{Y|A,X}(y | a, x) + s_X(x)]^2 P(dz) \right] \\ &\leq \frac{\sup_{y \in \mathcal{Y}} |\kappa(y, y)|}{\inf_{P' \in \mathcal{P}} \text{ess inf}_x g_{P'}(1 | x)} \left[\int [s_{Y|A,X}(y | a, x) + s_X(x)]^2 P(dz) \right] \end{aligned}$$

$$= \frac{\sup_{y \in \mathcal{Y}} |\kappa(y, y)|}{\inf_{P' \in \mathcal{P}} \text{ess inf}_x g_{P'}(1 \mid x)} \|s\|_{L^2(P)}^2, \quad (\text{S7})$$

where the first inequality holds by applying Jensen's inequality to bring the absolute value function inside the integral and then applying Cauchy-Schwarz to the (positive semidefinite) kernel κ , the second by Cauchy-Schwarz, and the third by Hölder's inequality with exponents $(p, q) = (1, \infty)$. The fraction on the right-hand side of (S7) is finite by the strong positivity assumption and the fact that κ is bounded, and so η_P is a bounded operator. Hence, by Lemma 2, ν is pathwise differentiable relative to \mathcal{P}_g with $\dot{\nu}_P = \eta_P$. In the same way as was done in (S4) for Example 1a, this pathwise differentiability over \mathcal{P}_g can be extended to show that ν is pathwise differentiable over the locally nonparametric model \mathcal{P} .

B.4.2 Efficient influence operator

Let $s \in \dot{\mathcal{P}}_P$ and $h \in \mathcal{H}$, and let $s_{Y|A,X}$ and s_X be as defined above (S6). We have that

$$\begin{aligned} \langle \dot{\nu}_P(s), h \rangle_{\mathcal{H}} &= \iint h(y) [s_{Y|A,X}(y \mid 1, x) + s_X(x)] P_{Y|A,X}(dy \mid 1, x) P_X(dx) \\ &= \int \frac{a}{g_P(1 \mid x)} \{h(y) - E[h(Y) \mid A = a, X = x]\} s(z) P(dz) \\ &\quad + \int (E_P[h(Y) \mid A = 1, X = x] - E_P E_P[h(Y) \mid A = 1, X]) s(z) P(dz). \end{aligned}$$

Hence,

$$\begin{aligned} \dot{\nu}_P^*(h)(z) &= \frac{a}{g_P(1 \mid x)} \{h(y) - E[h(Y) \mid A = 1, X = x]\} \\ &\quad + E_P[h(Y) \mid A = 1, X = x] - E_P E_P[h(Y) \mid A = 1, X]. \end{aligned}$$

B.4.3 Efficient influence function

By Theorem 1, the EIF will take the form

$$\begin{aligned} \phi_P(z)(y') &= \dot{\nu}_P^*(K_{y'})(z) = \frac{a}{g_P(1 \mid x)} \{\kappa(y, y') - E[\kappa(Y, y') \mid A = a, X = x]\} \\ &\quad + E_P[\kappa(Y, y') \mid A = 1, X = x] - E_P E_P[\kappa(Y, y') \mid A = 1, X]. \end{aligned}$$

provided we can show that this function belongs to $L^2(P; \mathcal{H})$. Defining $\mu_P^K(x) = E_P[K_Y \mid A = 1, X = x]$ and noting that $E_P \mu_P^K(x) = \nu(P)$, we can rewrite the above as follows:

$$\phi_P(z) = \frac{a}{g_P(1 \mid x)} [K_y - \mu_P^K(x)] + \mu_P^K(x) - \nu(P).$$

The fact that $\phi_P \in L^2(P; \mathcal{H})$ follows from the strong positivity assumption and the fact that the kernel κ is a bounded function.

B.4.4 Study of one-step estimator

The one-step estimator $\bar{\nu}_n := \frac{1}{2} \sum_{j=1}^2 [\nu(\hat{P}_n^j) + P_n^j \phi_n^j]$ that we study is a cross-fitted version of the estimator of the counterfactual kernel mean embedding introduced in Eq. 10 of Fawkes et al. (2022). Our general results provide several new results about this estimator that did not appear in that earlier work. First, Theorem 2 provides a set of conditions under which this estimator converges weakly to a tight limit. Second, when the conditions of Theorem 23 hold, (23) provides a precise sense in which $\bar{\nu}_n$ outperforms the inverse probability weighted estimator that was earlier introduced in Muandet et al. (2021). The earlier

work suggested that this estimator would be more efficient but did not provide any theoretical guarantees establishing this. Third, Theorems 3 and 4 provide a means to construct bootstrap-based confidence sets and hypothesis tests regarding the counterfactual mean embedding, with accompanying theoretical guarantees. Fawkes et al. (2022) proposed using an alternative, permutation-based procedure for making inference, but no theoretical guarantees were provided ensuring type I error control, consistency, or local power of the resulting test.

The calculations needed to establish that the conditions of our Theorem 2 hold are similar to those used to prove Theorem 1 in Fawkes et al. (2022) and those in our Appendix B.2.5, and therefore we only summarize the main findings here. For the remainder term $\mathcal{R}_P := \nu(P) + P\phi_0 - \nu(P_0)$, it holds that, for finite constants C_1 and C_2 that do not depend on $P \in \mathcal{P}$,

$$\begin{aligned} \|\mathcal{R}_P\|_{\mathcal{H}} &\leq C_1 \|g_P(1 \mid \cdot) - g_0(1 \mid \cdot)\|_{L^2(P_0, X)} \left[\int \|\mu_P^K - \mu_0^K\|_{\mathcal{H}}^2 P_{0, X}(dx) \right]^{1/2}, \\ \|\phi_P - \phi_0\|_{L^2(P_0, \mathcal{H})} &\leq C_2 \left(\|g_P(1 \mid \cdot) - g_0(1 \mid \cdot)\|_{L^2(P_0, X)} + \left[\int \|\mu_P^K - \mu_0^K\|_{\mathcal{H}}^2 P_{0, X}(dx) \right]^{1/2} \right). \end{aligned}$$

Taken together, these bounds show that the conditions of Theorem 2 will be satisfied in this example when, for $j \in \{1, 2\}$, \widehat{P}_n^j is such that $g_{\widehat{P}_n^j}$ and $\mu_{\widehat{P}_n^j}^K$ converge to g_0 and μ_0^K in probability according to the norms above and, moreover, the product of their rates of convergence is faster than $n^{-1/2}$.

B.5 Example 4: root-density function

Fix a distribution $P \in \mathcal{P}$, score s in the tangent set of \mathcal{P} at P , and submodel $\{P_\epsilon : \epsilon \in [0, \delta)\} \in \mathcal{P}(P, \mathcal{P}, s)$. Since ν is the square root of the density function in this example, the quadratic mean differentiability of $\{P_\epsilon : \epsilon \in [0, \delta)\}$ in (1) is, by definition, equivalent to the pathwise differentiability of ν as defined in (2), with the local parameter $\dot{\nu}_P(s)$ equal to $\epsilon s \nu(P)/2$. This local parameter is a bounded operator since, for any s ,

$$\|\dot{\nu}_P(s)\|_{L^2(\lambda)}^2 = \frac{1}{4} \int s^2 \nu(P)^2 d\lambda = \frac{1}{4} \|s\|_{L^2(P)}^2.$$

To verify the claimed form of the efficient influence operator given in Appendix A, we note that, for any $s \in \dot{\mathcal{P}}_P$ and $h \in L^2(\lambda)$,

$$\begin{aligned} \langle \dot{\nu}_P(s), h \rangle_{L^2(\lambda)} &= \frac{1}{2} \int h(z) s(z) \nu(P)(z) d\lambda(z) = \int \frac{h(z)}{2\nu(P)(z)} s(z) dP(z) \\ &= \int \left(\frac{h(z)}{2\nu(P)(z)} - E_P \left[\frac{h(Z)}{2\nu(P)(Z)} \right] \right) s(z) dP(z). \end{aligned}$$

As s and h were arbitrary, $\dot{\nu}_P^*(h)(z) = \frac{h(z)}{2\nu(P)(z)} - E_P \left[\frac{h(Z)}{2\nu(P)(Z)} \right]$.

B.6 Example 5: regression function

B.6.1 Pathwise differentiability

Fix a distribution $P \in \mathcal{P}$ and suppose that $\frac{d\lambda_X}{dP_X}$ is bounded P_X -almost surely. We prove that ν is pathwise differentiable at P relative to a locally nonparametric model and that $\dot{\nu}_P = \eta_P$, where

$$\eta_P(s)(x) = \int [y - \nu(P)(x)] s(x, y) P_{Y|X}(dy \mid x).$$

Fix a score function $s \in L_0^2(P)$. Let $\{P_\epsilon : \epsilon\} \in \mathcal{P}(P, \mathcal{P}, s)$ and $q_\epsilon := \sqrt{\frac{dP_\epsilon}{dP}}$. Let $q_{\epsilon, X} := \sqrt{\frac{dP_{\epsilon, X}}{dP_X}}$ denote the square root of the marginal density, $q_{\epsilon, Y|X}(\cdot | x) := \sqrt{\frac{dP_{\epsilon, Y|X}}{dP_{Y|X}}}(\cdot | x)$ the square root of the conditional density, $s_X(x) = E_P[s(X, Y) | X = x]$ and $s_{Y|X}(y | x) = s(x, y) - s_X(x)$. Observe that the following holds for P_X -almost all x :

$$\begin{aligned} & [\nu(P_\epsilon) - \nu(P) - \epsilon \eta_P(s)](x) \\ &= \int y [q_{\epsilon, Y|X}^2(y | x) - 1 - \epsilon s_{Y|X}(y | x)] P_{Y|X}(dy | x) \\ &= \int y \{[q_{\epsilon, Y|X}(y | x) - 1][q_{\epsilon, Y|X}(y | x) + 1] - \epsilon s_{Y|X}(y | x)\} P_{Y|X}(dy | x) \\ &= \int y \left\{ \left[q_{\epsilon, Y|X}(y | x) - 1 - \frac{\epsilon}{2} s_{Y|X}(y | x) \right] [q_{\epsilon, Y|X}(y | x) + 1] \right\} P_{Y|X}(dy | x) \\ &\quad + \frac{\epsilon}{2} \int y s_{Y|X}(y | x) [q_{\epsilon, Y|X}(y | x) - 1] P_{Y|X}(dy | x). \end{aligned}$$

For shorthand, we refer to the first term on the right as $A_\epsilon(x)$ and the second as $\frac{\epsilon}{2} B_\epsilon(x)$. We will show that $\|A_\epsilon\|_{L^2(\lambda_X)} = o(\epsilon)$ and $\|B_\epsilon\|_{L^2(\lambda_X)} = o(1)$. Combining this with the triangle inequality for the $L^2(\lambda_X)$ norm and the relation above will then give the result. For the first term, we note that the Cauchy-Schwarz inequality and the inequality $(a + b)^2 \leq 2a^2 + 2b^2$, we have, for P_X -almost all x ,

$$|A_\epsilon(x)|^2 \leq \left\| q_{\epsilon, Y|X} - 1 - \frac{\epsilon}{2} s_{Y|X} \right\|_{L^2(P_{Y|X=x})}^2 (2E_{P_\epsilon}[Y^2 | X = x] + 2E_P[Y^2 | X = x]).$$

Integrating both sides above against λ_X , applying Hölder's inequality with exponents $(1, \infty)$, and applying Lemma S8, we find that

$$\begin{aligned} \|A_\epsilon\|_{L^2(\lambda_X)}^2 &\leq 2 \operatorname{ess\,sup}_x \frac{d\lambda_X}{dP_X}(x) (E_{P_\epsilon}[Y^2 | X = x] + E_P[Y^2 | X = x]) \\ &\quad \cdot \left\| q_{\epsilon, Y|X} - 1 - \frac{\epsilon}{2} s_{Y|X} \right\|_{L^2(P)}^2 = o(\epsilon^2), \end{aligned}$$

where the essential supremum is over P_X . Above we used (S1) and the assumption that $\frac{d\lambda_X}{dP_X}$ is bounded with P_X -probability one.

We now show that $\|B_\epsilon\|_{L^2(\lambda_X)} = o(1)$. Let $B_{\epsilon,1}(x) = \int 1\{|ys_{Y|X}(y | x)| \leq \epsilon^{-1/2}\} y s_{Y|X}(y | x) [q_{\epsilon, Y|X}(y | x) - 1] P_{Y|X}(dy | x)$ and $B_{\epsilon,2}(x) = B_\epsilon(x) - B_{\epsilon,1}(x)$. By the triangle inequality, it suffices to show that $\|B_{\epsilon,j}\|_{L^2(\lambda_X)} = o(1)$, $j \in \{1, 2\}$. Using that $y^2 s_{Y|X}^2(y | x) 1\{|ys_{Y|X}(y | x)| \leq \epsilon^{-1/2}\} \leq \epsilon^{-1}$, Jensen's inequality, Hölder's inequality with exponents $(1, \infty)$, and Lemma S8,

$$\|B_{\epsilon,1}\|_{L^2(\lambda_X)}^2 \leq \epsilon^{-1} \left[\operatorname{ess\,sup}_x \frac{d\lambda_X}{dP_X}(x) \right] \|q_{\epsilon, Y|X} - 1\|_{L^2(P)}^2 = O(\epsilon),$$

where the essential supremum is over P_X . By the Cauchy-Schwartz inequality and the inequality $(a - b)^2 \leq 2(a^2 + b^2)$, the following holds for P_X -almost all x :

$$\begin{aligned} |B_{\epsilon,2}(x)|^2 &\leq 2(E_{P_\epsilon}[Y^2 | X = x] + E_P[Y^2 | X = x]) \\ &\quad \cdot \int s_{Y|X}^2(y | x) 1\{|ys_{Y|X}(y | x)| > \epsilon^{-1/2}\} P_{Y|X}(dy | x). \end{aligned}$$

Integrating both sides over λ_X and applying Hölder's inequality with exponents $(1, \infty)$ gives that

$$\|B_{\epsilon,2}\|_{L^2(\lambda_X)}^2 \leq 2 \left[\operatorname{ess\,sup}_x \frac{d\lambda_X}{dP_X}(x) (E_{P_\epsilon}[Y^2 | X = x] + E_P[Y^2 | X = x]) \right]$$

$$\begin{aligned} & \cdot \int s_{Y|X}^2(y|x) 1\{|ys_{Y|X}(y|x)| > \epsilon^{-1/2}\} dP(z) \\ & = o(1), \end{aligned}$$

where the essential supremum is over P_X . Thus, $\|\nu(P_\epsilon) - \nu(P) - \epsilon\eta_P(s)\|_{L^2(\lambda_X)} = o(\epsilon)$.

We now prove that η_P is a bounded operator. For any $s \in \dot{\mathcal{P}}_P$, applying Cauchy-Schwarz followed by Hölder's inequality with exponents $(1, \infty)$ shows that

$$\begin{aligned} \|\eta_P(s)\|_{L^2(\lambda_X)}^2 &= \int \frac{d\lambda_X}{dP_X}(x) \left[\int [y - \nu(P)(x)] s_{Y|X}(y|x) P_{Y|X}(dy|x) \right]^2 P_X(dx) \\ &\leq \int \frac{d\lambda_X}{dP_X}(x) \text{Var}_P(Y|X=x) \left[\int |s_{Y|X}(y|x)|^2 P_{Y|X}(dy|x) \right] P_X(dx) \\ &\leq \left[\text{ess sup}_x \frac{d\lambda_X}{dP_X}(x) \text{Var}_P(Y|X=x) \right] \|s\|_{L^2(P)}^2, \end{aligned}$$

where the essential supremum is over P_X . Hence,

$$\|\eta_P\|_{\text{op}} \leq \text{ess sup}_x \sqrt{\frac{d\lambda_X}{dP_X}(x) \text{Var}_P(Y|X=x)},$$

which is finite by (S1) and the assumption that $\frac{d\lambda_X}{dP_X}$ is P_X -a.s. bounded. Since η_P is also linear, ν is pathwise differentiable with local parameter $\dot{\nu}_P = \eta_P$.

B.6.2 Efficient influence operator

For any $h \in L^2(\lambda_X)$ and $s \in \dot{\mathcal{P}}_P$,

$$\begin{aligned} \langle \dot{\nu}_P(s), h \rangle_{L^2(\lambda_X)} &= \iint [y - \nu(P)(x)] h(x) s(x, y) P_{Y|X}(dy|x) \lambda_X(dx) \\ &= \iint \frac{d\lambda_X}{dP_X}(x) [y - \nu(P)(x)] h(x) s(x, y) P_{Y|X}(dy|x) P_X(dx) \\ &= \int \frac{d\lambda_X}{dP_X}(x) [y - \nu(P)(x)] h(x) s(x, y) P(dz) = \langle \dot{\nu}_P^*(h), s \rangle_{L^2(P)}, \end{aligned}$$

where $\dot{\nu}_P^*(h)(z) = \frac{d\lambda_X}{dP_X}(x) [y - \nu(P)(x)] h(x)$. Hence, $\dot{\nu}_P^*$ is the efficient influence operator.

B.7 Example 6: kernel mean embedding

The parameter considered in this example is a special case of the counterfactual kernel mean embedding parameter considered in Example 3 when $A = 1$ almost surely. Consequently, the proof of the pathwise differentiability of the parameter in this example, and also the calculation of its efficient influence operator and EIF, follow directly from those in Appendix B.4.

B.8 Example 7: conditional average treatment effect

The proof we provide does not require a new application of Lemma 2, but instead leverages the application of that lemma that we already worked out in Example 5. First, we establish that ν is pathwise differentiable with the claimed local parameter relative to a semiparametric model where the propensity is known. We do this by leveraging the result from Example 5 to establish the pathwise differentiability of a regression of a certain pseudo-outcome against the covariates. Second, we establish that working in the larger, locally nonparametric model where this quantity is not known does not change this result: ν is still pathwise differentiable with the same local parameter. An alternative argument, which we do not give here, would

entail directly applying Lemma 2 when establishing pathwise differentiability in the semiparametric model considered in the first step.

Fix a distribution P_0 in the locally nonparametric model \mathcal{P} that is such that $\frac{d\lambda_X}{dP_{0,X}}$ is bounded $P_{0,X}$ -a.s. and let $g := g_{P_0}$. We will establish pathwise differentiability at P_0 with the claimed efficient influence operator, and, since P_0 was arbitrary, this will establish the desired result. For a distribution $P \in \mathcal{P}$, let P_g denote the distribution of Z that has the same conditional distribution of $Y | A, X$ and marginal distribution of X as P , but has propensity equal to g ; in other words, for all bounded, continuous functions $f : \mathcal{Z} \rightarrow \mathbb{R}$, $E_{P_g}[f(Z)] = \int \sum_{a=0}^1 \int f(x, a, y) P_{Y|A,X}(dy | a, x) g(a | x) P_X(dx)$. The semiparametric model that we study is given by $\mathcal{P}_g := \{P_g : P \in \mathcal{P}\}$. Without loss of generality, we suppose that $\mathcal{P}_g \subseteq \mathcal{P}$; if this is not the case, then we can simply extend the definition of ν to \mathcal{P}_g by letting $\nu(P_g) = \nu(P)$ for any $P \in \mathcal{P}$. For any $P_g \in \mathcal{P}_g$, it can be verified that, for λ_X -almost all X ,

$$\nu(P_g)(x) = E_{P_g}[W | X = x],$$

where $W := (2A - 1)Y/g(A | X)$. The above suggests that we can use our results from the regression setting in Example 5 to derive the efficient influence operator of ν relative to \mathcal{P}_g . To this end, we define the model $\tilde{\mathcal{P}}_g := \{P_g \circ f_g^{-1} : P_g \in \mathcal{P}_g\}$, where $P_g \circ f_g^{-1}$ is the pushforward measure of P_g under $f_g(x, a, y) := (x, w) := (x, [2a - 1]y/g(a | x))$. We also define the parameter $\tilde{\nu} : \tilde{\mathcal{P}}_g \rightarrow \mathcal{H}$ so that $\tilde{\nu}(P_g \circ f_g^{-1}) = \nu(P_g)$, where this definition is valid even if there are two distinct distributions $P_g, P'_g \in \mathcal{P}_g$ that make it so that $P_g \circ f_g^{-1} = P'_g \circ f_g^{-1}$ since the preceding display shows that $\nu(P_g) = \nu(P'_g)$ in this case. The parameter $\tilde{\nu}$ takes as input a distribution of features X and an outcome W from a locally nonparametric model and outputs a regression function. Consequently, the results of Example 5 imply that, for any P_g such that $\frac{dP_{g,X}}{d\lambda_X}$ is bounded, this parameter is pathwise differentiable at $\tilde{P}_g := P_g \circ f_g^{-1}$ with local parameter

$$\dot{\tilde{\nu}}_{\tilde{P}_g}(\tilde{s})(x) = \int [w - \tilde{\nu}(\tilde{P}_g)(x)] \tilde{s}(x, w) \tilde{P}_g(dw | x).$$

We now use the pathwise differentiability of $\tilde{\nu}$ relative to $\tilde{\mathcal{P}}_g$ to establish the pathwise differentiability of ν relative to \mathcal{P}_g . To this end, let $\{P_{g,\epsilon} : \epsilon \in [0, \delta)\} \in \mathcal{P}(P_g, \mathcal{P}_g, s_g)$, where $P_g \in \mathcal{P}_g$ and s_g belong to the tangent set of \mathcal{P}_g at P_g . Similar arguments to those used to establish Lemma S8 can be used to show that $\{P_{g,\epsilon} \circ f_g^{-1} : \epsilon \in [0, \delta)\} \in \mathcal{P}(\tilde{P}_g, \tilde{\mathcal{P}}_g, \tilde{s}_g)$, where $\tilde{s}_g(x, w) := E_{P_g}[s_g(X, A, Y) | X = x, W = w]$. Combining this with the pathwise differentiability of $\tilde{\nu}$ relative to $\tilde{\mathcal{P}}_g$ and the definition of $\tilde{\nu}$ shows that

$$\begin{aligned} & \left\| \nu(P_{g,\epsilon}) - \nu(P_g) - \epsilon \dot{\tilde{\nu}}_{P_g \circ f_g^{-1}}(\tilde{s}_g) \right\|_{\mathcal{H}} \\ &= \left\| \tilde{\nu}(P_{g,\epsilon} \circ f_g^{-1}) - \tilde{\nu}(P_g \circ f_g^{-1}) - \epsilon \dot{\tilde{\nu}}_{P_g \circ f_g^{-1}}(\tilde{s}_g) \right\|_{\mathcal{H}} = o(\epsilon). \end{aligned}$$

Since the operator $\dot{\nu}_{P_g} : \dot{\mathcal{P}}_{g,P_g} \rightarrow \mathcal{H}$ defined by $\dot{\nu}_{P_g}(s_g) := \dot{\nu}_{\tilde{P}_g}(\tilde{s}_g)$ is bounded and linear, where $\dot{\mathcal{P}}_{g,P_g}$ denotes the tangent space of \mathcal{P}_g at P_g , this shows that ν is pathwise differentiable at P_g relative to \mathcal{P}_g with local parameter $\dot{\nu}_{P_g}$.

We now use the pathwise differentiability of ν relative to \mathcal{P}_g to establish its pathwise differentiability at P_0 relative to \mathcal{P} . Let $\{P_\epsilon : \epsilon \in [0, \delta)\} \in \mathcal{P}(P_0, \mathcal{P}, s)$, where s is the tangent set of \mathcal{P} at P_0 . Letting $P_{g,\epsilon}$ be the distribution that has the same conditional distribution of $Y | A, X$ and marginal distribution of X as under P_ϵ but with propensity g , it can be shown that $\{P_{g,\epsilon} : \epsilon \in [0, \delta)\} \in \mathcal{P}(P_0, \mathcal{P}_g, s_g)$, where $s_g(z) = s(z) - E_0[s(Z) | A = a, X = x] + E_0[s(Z) | X = x]$. Combining this with the facts that $P_{g,0} = P_0$ and ν is invariant to changes in the propensity of its input, we find that

$$\|\nu(P_\epsilon) - \nu(P_0) - \epsilon \dot{\nu}_0(s_g)\|_{\mathcal{H}} = \|\nu(P_{g,\epsilon}) - \nu(P_g) - \epsilon \dot{\nu}_0(s_g)\|_{\mathcal{H}} = o(\epsilon).$$

The above establishes that ν is pathwise differentiable at P_0 relative to \mathcal{P} , with local parameter

$$\begin{aligned}
s \mapsto \dot{\nu}_0(s_g) &= \dot{\tilde{\nu}}_{P_0 \circ f_g^{-1}}(\tilde{s}_g) \\
&= \int [w - \nu(P_0)(x)] E_0[s_g(X, A, Y) \mid X = x, W = w] (P_0 \circ f_g^{-1})(dw \mid x) \\
&= E_0 \{ [W - \nu(P_0)(X)] E_0[s_g(X, A, Y) \mid X, W] \mid X = x \} \\
&= E_0 \{ [W - \nu(P_0)(X)] s_g(X, A, Y) \mid X = x \} \\
&= E_0 \{ [W - E_0(W \mid A, X) + E_0(W \mid X) - \nu(P_0)(X)] s_g(X, A, Y) \mid X = x \} \\
&= E_0 \{ [W - E_0(W \mid A, X)] s_g(X, A, Y) \mid X = x \} \\
&= E_0 \{ [W - E_0(W \mid A, X)] s(X, A, Y) \mid X = x \} \\
&= E_0 \left\{ \left[\frac{2A-1}{g_{P_0}(A \mid X)} \{Y - \mu_{P_0, A}(X)\} \right] s(X, A, Y) \mid X = x \right\},
\end{aligned}$$

which matches the claimed form of the local parameter from (S2). To verify that the efficient influence operator takes the form in (S3), it can be directly established that $\langle \dot{\nu}_0^*(h), s \rangle_{L^2(P_0)} = \langle h, \dot{\nu}_0(s) \rangle_{\mathcal{H}}$ for all $h \in \mathcal{H}$ and $s \in L_0^2(P_0)$. The calculations to establish this are straightforward and so are omitted.

C Proofs of results from the main text, and supporting lemmas

C.1 Proofs for Section 2

C.1.1 Proofs for Section 2.3

Lemma S2. Suppose $\dot{\mathcal{H}}_P$ is an RKHS, $\nu : \mathcal{P} \rightarrow \mathcal{H}$ is pathwise differentiable at P , and $\tilde{\phi}_P$ as defined in (5) is P -Bochner square integrable. For all $h \in \mathcal{H}$, define $\langle \tilde{\phi}_P, h \rangle_{\mathcal{H}} : \mathcal{Z} \rightarrow \mathbb{R}$ so that $\langle \tilde{\phi}_P, h \rangle_{\mathcal{H}}(z) = \langle \tilde{\phi}_P(z), h \rangle_{\mathcal{H}}$. Then, $\langle \tilde{\phi}_P, h \rangle_{\mathcal{H}} \in \dot{\mathcal{P}}_P$ and $\langle \tilde{\phi}_P, h \rangle_{\mathcal{H}} = \dot{\nu}_P^*(h)$ P -almost surely.

Proof. Fix $h \in \mathcal{H}$. The fact that $\langle \tilde{\phi}_P, h \rangle_{\mathcal{H}} \in L^2(P)$ follows from Cauchy-Schwarz and the fact that $\tilde{\phi}_P \in L^2(P; \mathcal{H})$. In particular,

$$\int |\langle \tilde{\phi}_P(z), h \rangle_{\mathcal{H}}|^2 P(dz) \leq \|h\|_{\mathcal{H}}^2 \|\tilde{\phi}_P\|_{L^2(P; \mathcal{H})}^2 < \infty.$$

Let s^\perp be an element of the orthogonal complement of the tangent space $\dot{\mathcal{P}}_P$. Then,

$$\begin{aligned}
\int \langle \tilde{\phi}_P(z), h \rangle_{\mathcal{H}} s^\perp(z) P(dz) &= \int \langle t \mapsto \dot{\nu}_P^*(K_t)(z), h \rangle_{\mathcal{H}} s^\perp(z) P(dz) \\
&= \left\langle t \mapsto \int \dot{\nu}_P^*(K_t)(z) s^\perp(z) P(dz), h \right\rangle_{\mathcal{H}} \\
&= \langle t \mapsto \langle \dot{\nu}_P^*(K_t), s^\perp \rangle_{L^2(P)}, h \rangle_{\mathcal{H}} = 0,
\end{aligned}$$

where we use the P -Bochner square integrability of $\tilde{\phi}_P$ to interchange the integral and the inner product and use that $\dot{\nu}_P^*(K_t) \in \dot{\mathcal{P}}_P$ for all $t \in \mathcal{T}$. Note also that $\langle \dot{\nu}_P^*(h), s^\perp \rangle_{L^2(P)} = 0$ since $\dot{\nu}_P^*(h) \in \dot{\mathcal{P}}_P$. Hence, for any s^\perp in the orthogonal complement of $\dot{\mathcal{P}}_P$,

$$\left\langle \langle \tilde{\phi}_P, h \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h), s^\perp \right\rangle_{L^2(P)} = 0. \quad (\text{S8})$$

Let $s \in \dot{\mathcal{P}}_P$. Then, following some of the same calculations as earlier,

$$\int \langle \tilde{\phi}_P(z), h \rangle_{\mathcal{H}} s(z) P(dz) = \langle t \mapsto \langle \dot{\nu}_P^*(K_t), s \rangle_{L^2(P)}, h \rangle_{\mathcal{H}}.$$

Furthermore, for any $t \in \mathcal{T}$, $\langle \dot{\nu}_P^*(K_t), s \rangle_{L^2(P)} = \langle K_t, \dot{\nu}_P(s) \rangle_{\mathcal{H}}$. Since $\dot{\nu}_P(s)$ belongs to the RKHS $\dot{\mathcal{H}}_P$, $\langle K_t, \dot{\nu}_P(s) \rangle_{\mathcal{H}} = \dot{\nu}_P(s)(t)$. Plugging these observations into the above shows that the right-hand side is equal to $\langle \dot{\nu}_P(s), h \rangle_{\mathcal{H}} = \langle \dot{\nu}_P^*(h), s \rangle_{L^2(P)}$, and so

$$\left\langle \langle \tilde{\phi}_P, h \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h), s \right\rangle_{L^2(P)} = 0.$$

As the above holds for all s in the tangent space $\dot{\mathcal{P}}_P$ and (S8) holds for all s^\perp in its orthogonal complement, $\langle \tilde{\phi}_P, h \rangle_{\mathcal{H}} = \dot{\nu}_P^*(h)$ P -almost surely. \square

Lemma S3. *In the setting of Lemma S2, it is P -a.s. true that $\sup_{h \in \mathcal{H}} |\langle \tilde{\phi}_P, h \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h)| = 0$. Hence, for P -almost all z , $\dot{\nu}_P^*(\cdot)(z) : \mathcal{H} \rightarrow \mathbb{R}$ is a bounded linear functional with Riesz representation $\tilde{\phi}_P(z)$. In other words, $\tilde{\phi}_P$ is the EIF of ν at P .*

Proof. Since we have assumed throughout that a separable version of the efficient influence process is used, there exists a countable dense subset \mathcal{H}' of \mathcal{H} and a P -probability one subset \mathcal{Z}' of \mathcal{Z} such that, for all $h \in \mathcal{H}$ and $z \in \mathcal{Z}'$, there exists an \mathcal{H}' -valued sequence $(h'_j)_{j=1}^\infty$ that converges to h and satisfies $\dot{\nu}_P^*(h'_j)(z) \rightarrow \dot{\nu}_P^*(h)(z)$ as $j \rightarrow \infty$. For each $h \in \mathcal{H}'$, we let

$$\mathcal{Z}_h'' = \{z \in \mathcal{Z}' : \langle \tilde{\phi}_P(z), h \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h)(z) = 0\},$$

and we define $\mathcal{Z}'' := \bigcap_{h \in \mathcal{H}'} \mathcal{Z}_h''$. By Lemma S2 and the fact that \mathcal{Z}' is a P -probability one set, $P(\mathcal{Z}_h'') = 1$ for each $h \in \mathcal{H}'$ and, as \mathcal{H}' is countable, $P(\mathcal{Z}'') = 1$ as well. In what follows we will show that, for all $z \in \mathcal{Z}''$, $\sup_{h \in \mathcal{H}} |\langle \tilde{\phi}_P(z), h \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h)(z)| = 0$. To this end, fix $z \in \mathcal{Z}''$ and $\epsilon > 0$ and let $h_\epsilon \in \mathcal{H}$ be such that

$$|\langle \tilde{\phi}_P(z), h_\epsilon \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h_\epsilon)(z)| \geq \sup_{h \in \mathcal{H}} |\langle \tilde{\phi}_P(z), h \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h)(z)| - \epsilon.$$

By the separability of the efficient influence process, there exists an \mathcal{H}' -valued sequence $(h_{\epsilon,j})$ that converges to h_ϵ that is such that $\dot{\nu}_P^*(h_{\epsilon,j})(z) \rightarrow \dot{\nu}_P^*(h_\epsilon)(z)$ as $j \rightarrow \infty$. By the continuity of the inner product $\langle \tilde{\phi}_P(z), \cdot \rangle_{\mathcal{H}}$, it also holds that $\langle \tilde{\phi}_P(z), h_{\epsilon,j} \rangle_{\mathcal{H}} \rightarrow \langle \tilde{\phi}_P(z), h_\epsilon \rangle_{\mathcal{H}}$ as $j \rightarrow \infty$. Consequently, $|\langle \tilde{\phi}_P(z), h_{\epsilon,j} \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h_{\epsilon,j})(z)|$ converges to the left-hand side above as $j \rightarrow \infty$, and so there exists a sufficiently large j such that

$$|\langle \tilde{\phi}_P(z), h_{\epsilon,j} \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h_{\epsilon,j})(z)| \geq \sup_{h \in \mathcal{H}} |\langle \tilde{\phi}_P(z), h \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h)(z)| - 2\epsilon.$$

As z is in \mathcal{Z}'' , z is in $\mathcal{Z}_{h_{\epsilon,j}}''$ as well. Hence, the left-hand side above is zero, which shows that

$$\sup_{h \in \mathcal{H}} |\langle \tilde{\phi}_P(z), h \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h)(z)| \leq 2\epsilon.$$

As $\epsilon > 0$ was arbitrary, the left-hand side above is equal to zero. This proves the first claim of the lemma. The second claim follows directly from the fact that $h \mapsto \langle \tilde{\phi}_P(z), h \rangle$ is a bounded linear functional and by the definition of the Riesz representation of such a functional. \square

Lemma S4. *Let $\nu : \mathcal{P} \rightarrow \mathcal{H}$ be pathwise differentiable at P and suppose that $\dot{\mathcal{H}}_P$ is an RKHS. If ν has EIF ϕ_P at P , then $\phi_P = \tilde{\phi}_P$ P -a.s., where $\tilde{\phi}_P$ is as defined in (5).*

Proof. Since ϕ_P is the EIF of ν at P , there exists a P -probability-one set \mathcal{Z}' such that, for all $z \in \mathcal{Z}'$, $\dot{\nu}_P^*(\cdot)(z) : \mathcal{H} \rightarrow \mathbb{R}$ is a bounded linear functional with Riesz representation ϕ_P . In other words, for all $z \in \mathcal{Z}'$,

$$\sup_{h \in \mathcal{H}} |\langle \phi_P(z), h \rangle_{\mathcal{H}} - \dot{\nu}_P^*(h)(z)| = 0.$$

Fix $z \in \mathcal{Z}'$. Since $\dot{\mathcal{H}}_P \subseteq \mathcal{H}$, the above shows that $\langle \phi_P(z), K_t \rangle_{\mathcal{H}} - \dot{\nu}_P^*(K_t)(z) = 0$ for all $t \in \mathcal{T}$. Since $\phi_P(z) \in \dot{\mathcal{H}}_P$, $\phi_P(z)(t) = \langle \phi_P(z), K_t \rangle_{\mathcal{H}}$ for all $t \in \mathcal{T}$. Combining this with the fact that $\tilde{\phi}_P(t) := \dot{\nu}_P^*(K_t)(z)$, this shows that $\phi_P(z) = \tilde{\phi}_P(z)$. As z is an arbitrary element of the P -probability one set \mathcal{Z}' , this shows that $\phi_P = \tilde{\phi}_P$ P -almost surely. \square

Proof of Theorem 1. The first statement, (i), was established in Lemma S4. The second statement, (ii), was established in Lemma S3. \square

Lemma S5. *Let \mathcal{P} be a statistical model of distributions that are equivalent in that, for all $P_1, P_2 \in \mathcal{P}$, $P_1 \ll P_2$ and $P_2 \ll P_1$. Let $\{P_{\epsilon} : \epsilon \in [0, \delta]\} \in \mathcal{P}(P, \mathcal{P}, s)$ be a quadratic mean differentiable submodel. Let $\nu : \mathcal{P} \rightarrow \mathcal{H}$ be pathwise differentiable at $P \in \mathcal{P}$ with a P -almost surely bounded EIF ϕ_P , in the sense that $\|\phi_P(Z)\|_{\mathcal{H}}$ is a bounded random variable when $Z \sim P$. Under these conditions,*

$$(P_{\epsilon} - P)\phi_P - \epsilon \dot{\nu}_P(s) = o(\epsilon). \quad (\text{S9})$$

The above lemma requires that $\|\phi_P(Z)\|_{\mathcal{H}}$ be a bounded random variable in order to show that (S9) holds. Lemma S6 will provide an alternative condition under which (S9) holds. In particular, rather than impose a boundedness condition on the EIF, that lemma will require that the submodel be approximately linear, in the sense that $\frac{dP_{\epsilon}}{dP} \approx 1 + \epsilon s$ in an appropriate sense.

Proof of Lemma S5. Fix a quadratic mean differentiable submodel $\{P_{\epsilon} : \epsilon \in [0, \delta]\} \in \mathcal{P}(P, \mathcal{P}, s)$. For $\epsilon \in [0, \delta]$, let $h_{\epsilon} := (P_{\epsilon} - P)\phi_P - \epsilon \dot{\nu}_P(s)$. We will show that $\|h_{\epsilon}\|_{\mathcal{H}} = o(\epsilon)$. Let $g_{\epsilon} := h_{\epsilon}/\|h_{\epsilon}\|_{\mathcal{H}}$, where we use the convention that $g_{\epsilon} = 0$ when $h_{\epsilon} = 0$. Observe that

$$\|h_{\epsilon}\|_{\mathcal{H}} = \langle (P_{\epsilon} - P)\phi_P - \epsilon \dot{\nu}_P(s), h_{\epsilon} \rangle_{\mathcal{H}} = \langle (P_{\epsilon} - P)\phi_P, g_{\epsilon} \rangle_{\mathcal{H}} - \epsilon \langle s, \dot{\nu}_P^*(g_{\epsilon}) \rangle_{L^2(P)}. \quad (\text{S10})$$

We now study the inner product $\langle (P_{\epsilon} - P)\phi_P, h_{\epsilon} \rangle_{\mathcal{H}}$ that appears above. Since $\phi_P(Z)$ is bounded under sampling from P and $P_{\epsilon} \ll P$, $\phi_P : \mathcal{Z} \rightarrow \mathcal{H}$ is Bochner integrable both under sampling from P and P_{ϵ} . Consequently, $\langle (P_{\epsilon} - P)\phi_P, g_{\epsilon} \rangle_{\mathcal{H}} = \int \langle \phi_P(z), g_{\epsilon}(z) \rangle_{\mathcal{H}} d(P_{\epsilon} - P)(z)$. Adding and subtracting terms from this identity and letting $q_{\epsilon} := p_{\epsilon}^{1/2}$ and $q := p^{1/2}$ yields that

$$\begin{aligned} & \langle (P_{\epsilon} - P)\phi_P, g_{\epsilon} \rangle_{\mathcal{H}} \\ &= \int \langle \phi_P(z), g_{\epsilon}(z) \rangle_{\mathcal{H}} [q_{\epsilon}(z) + q(z)][q_{\epsilon}(z) - q(z)] d\lambda(z) \\ &= \epsilon \int \langle \phi_P(z), g_{\epsilon}(z) \rangle_{\mathcal{H}} s(z) q^2(z) d\lambda(z) + \frac{1}{2} \epsilon \int \langle \phi_P(z), g_{\epsilon}(z) \rangle_{\mathcal{H}} [q_{\epsilon}(z) - q(z)] s(z) q(z) d\lambda(z) \\ & \quad + \int \langle \phi_P(z), g_{\epsilon}(z) \rangle_{\mathcal{H}} [q_{\epsilon}(z) + q(z)] \left[q_{\epsilon}(z) - q(z) - \frac{1}{2} \epsilon s(z) q(z) \right] d\lambda(z). \end{aligned}$$

By the definition of the EIF, $\langle \phi_P(z), g_{\epsilon}(z) \rangle_{\mathcal{H}} = \dot{\nu}_P^*(g_{\epsilon})(z)$ P -almost surely. Hence, the first term on the right-hand side above is equal to $\epsilon \langle s, \dot{\nu}_P^*(g_{\epsilon}) \rangle_{L^2(P)}$, and so (S10) shows that

$$\|h_{\epsilon}\|_{\mathcal{H}} = \frac{1}{2} \epsilon \int \langle \phi_P(z), g_{\epsilon}(z) \rangle_{\mathcal{H}} [q_{\epsilon}(z) - q(z)] s(z) q(z) d\lambda(z)$$

$$+ \int \langle \phi_P(z), g_\epsilon \rangle_{\mathcal{H}} [q_\epsilon(z) + q(z)] \left[q_\epsilon(z) - q(z) - \frac{1}{2} \epsilon s(z) q(z) \right] d\lambda(z).$$

By Jensen's inequality and Cauchy-Schwarz, this yields that

$$\begin{aligned} \|h_\epsilon\|_{\mathcal{H}} &\leq \frac{1}{2} \epsilon \|g_\epsilon\|_{\mathcal{H}} \int \|\phi_P(z)\|_{\mathcal{H}} |q_\epsilon(z) - q(z)| s(z) q(z) d\lambda(z) \\ &\quad + \|g_\epsilon\|_{\mathcal{H}} \int \|\phi_P(z)\|_{\mathcal{H}} [q_\epsilon(z) + q(z)] \left| q_\epsilon(z) - q(z) - \frac{1}{2} \epsilon s(z) q(z) \right| d\lambda(z). \end{aligned} \quad (\text{S11})$$

Since $\|g_\epsilon\|_{\mathcal{H}}$ is either equal to 1, if $h_\epsilon \neq 0$, or is equal to zero otherwise, we can establish that $\|h_\epsilon\|_{\mathcal{H}} = o(\epsilon)$ by showing that the first integral above is $o(1)$ and the second is $o(\epsilon)$. To show that the first integral is $o(1)$, we use (i) Cauchy-Schwarz, (ii) the fact that ϕ_P is essentially bounded, (iii) $s \in L^2(P)$, and (iv) the fact that the quadratic mean differentiability of $\{P_\epsilon : \epsilon \in [0, \delta)\}$ implies that $\|q_\epsilon - q\|_{L^2(\lambda)} = O(\epsilon)$. Combining these observations yields the display

$$\begin{aligned} &\int \|\phi_P(z)\|_{\mathcal{H}} |q_\epsilon(z) - q(z)| s(z) q(z) d\lambda(z) \\ &\leq \|\|\phi_P(\cdot)\|_{\mathcal{H}} s(\cdot) q(\cdot)\|_{L^2(\lambda)} \|q_\epsilon - q\|_{L^2(\lambda)} = O(\epsilon) = o(1). \end{aligned}$$

For the second integral in (S11), (i), (ii), and (iii), together with the inequality $(a + b)^2 \leq 2(a^2 + b^2)$ and the quadratic mean differentiability of $\{P_\epsilon : \epsilon \in [0, \delta)\}$, yield that

$$\begin{aligned} &\int \|\phi_P(z)\|_{\mathcal{H}} [q_\epsilon(z) + q(z)] \left| q_\epsilon(z) - q(z) - \frac{1}{2} \epsilon s(z) q(z) \right| d\lambda(z) \\ &\leq \|\|\phi_P(\cdot)\|_{\mathcal{H}} (q_\epsilon + q)\|_{L^2(\lambda)} \left\| q_\epsilon - q - \frac{1}{2} \epsilon s q \right\|_{L^2(\lambda)} \\ &\leq 2^{1/2} \left[\int \|\phi_P(z)\|_{\mathcal{H}}^2 d(P_\epsilon + P)(z) \right]^{1/2} \|q_\epsilon + q\|_{L^2(\lambda)} \left\| q_\epsilon - q - \frac{1}{2} \epsilon s q \right\|_{L^2(\lambda)} \\ &= o(\epsilon). \end{aligned}$$

Plugging the preceding two displays into (S11) completes the proof. \square

Lemma S6. *Fix a score s in the tangent set of \mathcal{P} at P and let $\{P_\epsilon : \epsilon \in [0, \delta)\} \in \mathcal{P}(P, \mathcal{P}, s)$ be such that*

$$\|dP_\epsilon/dP - 1 - \epsilon s\|_{L^2(P)} = o(\epsilon). \quad (\text{S12})$$

If ν is pathwise differentiable at P , then

$$\sup_{h \in \mathcal{H}_1} |P_\epsilon \dot{\nu}_P^*(h) + \langle \nu(P) - \nu(P_\epsilon), h \rangle_{\mathcal{H}}| = \sup_{h \in \mathcal{H}_1} |P_\epsilon \dot{\nu}_P^*(h) + \epsilon \langle \dot{\nu}_P(s), h \rangle_{\mathcal{H}}| + o(\epsilon) = o(\epsilon), \quad (\text{S13})$$

where \mathcal{H}_1 denotes the unit ball in \mathcal{H} . Moreover, if ν has P -Bochner square integrable EIF ϕ_P at P , then

$$(P_\epsilon - P)\phi_P + \epsilon \dot{\nu}_P(s) = o(\epsilon).$$

When s is bounded and \mathcal{P} is nonparametric, there is a quadratic mean differentiable submodel $\{P_\epsilon : \epsilon \in [0, \delta)\} \in \mathcal{P}(P, \mathcal{P}, s)$ that is such that $\frac{dP_\epsilon}{dP} = 1 + \epsilon s$. The condition in (S12) holds trivially for this submodel, since $\|dP_\epsilon/dP - 1 - \epsilon s\|_{L^2(P)} = 0$. This condition will generally also hold for many other quadratic mean differentiable submodels.

Proof of Lemma S6. Suppose that s and $\{P_\epsilon : \epsilon \in [0, \delta)\}$ are as in the statement of the lemma and that

ν is pathwise differentiable at P . For any $h \in \mathcal{H}$, the fact that $P\dot{\nu}_P^*(g) = 0$ implies that

$$P_\epsilon\dot{\nu}_P^*(h) = \epsilon\langle s, \dot{\nu}_P^*(h) \rangle_{L^2(P)} + P \left[\left(\frac{dP_\epsilon}{dP} - 1 - \epsilon s \right) \dot{\nu}_P^*(h) \right].$$

Combining this with the fact that $\langle s, \dot{\nu}_P^*(h) \rangle_{L^2(P)} = \langle \dot{\nu}_P(s), h \rangle_{\mathcal{H}}$ shows that

$$\begin{aligned} P_\epsilon\dot{\nu}_P^*(h) + \langle \nu(P) - \nu(P_\epsilon), h \rangle_{\mathcal{H}} \\ = \langle \nu(P) - \nu(P_\epsilon) - \epsilon\dot{\nu}_P(s), h \rangle_{\mathcal{H}} + P \left[\left(\frac{dP_\epsilon}{dP} - 1 - \epsilon s \right) \dot{\nu}_P^*(h) \right]. \end{aligned}$$

Taking an absolute value and then a supremum over $h \in \mathcal{H}_1$ and subsequently applying the triangle inequality and Cauchy-Schwarz yields that

$$\begin{aligned} \sup_{h \in \mathcal{H}_1} |P_\epsilon\dot{\nu}_P^*(h) + \langle \nu(P) - \nu(P_\epsilon), h \rangle_{\mathcal{H}}| \\ \leq \sup_{h \in \mathcal{H}_1} |\langle \nu(P) - \nu(P_\epsilon) - \epsilon\dot{\nu}_P(s), h \rangle_{\mathcal{H}}| + \sup_{h \in \mathcal{H}_1} \left| P \left[\left(\frac{dP_\epsilon}{dP} - 1 - \epsilon s \right) \dot{\nu}_P^*(h) \right] \right| \\ \leq \|\nu(P) - \nu(P_\epsilon) - \epsilon\dot{\nu}_P(s)\|_{\mathcal{H}} + \left\| \frac{dP_\epsilon}{dP} - 1 - \epsilon s \right\|_{L^2(P)} \|\dot{\nu}_P^*\|_{\text{op}} \end{aligned}$$

The first term on the right-hand side is $o(\epsilon)$ by the pathwise differentiability of ν , and the second is $o(\epsilon)$ by (S12) and the fact that $\dot{\nu}_P^*$ is a bounded operator. Eq. S13 follows by combining the above with the fact that, by the pathwise differentiability of ν ,

$$\sup_{h \in \mathcal{H}_1} |P_\epsilon\dot{\nu}_P^*(h) + \langle \nu(P) - \nu(P_\epsilon), h \rangle_{\mathcal{H}}| = \sup_{h \in \mathcal{H}_1} |P_\epsilon\dot{\nu}_P^*(h) + \epsilon\langle \dot{\nu}_P(s), h \rangle_{\mathcal{H}}| + o(\epsilon).$$

Now suppose that ν has a P -Bochner square integrable EIF ϕ_P at P . We have that

$$\begin{aligned} & \| (P_\epsilon - P)\phi_P + \epsilon\dot{\nu}_P(s) \|_{\mathcal{H}} \\ &= \| (P_\epsilon - P)\phi_P + \nu(P_\epsilon) - \nu(P) \|_{\mathcal{H}} + o(\epsilon) \\ &= \sup_{h \in \mathcal{H}_1} [\langle P_\epsilon\phi_P + \nu(P_\epsilon) - \nu(P), h \rangle_{\mathcal{H}}] + o(\epsilon) \\ &\leq \sup_{h \in \mathcal{H}_1} |\langle P_\epsilon\phi_P, h \rangle_{\mathcal{H}} - P_\epsilon\dot{\nu}_P^*(h)| + \sup_{h \in \mathcal{H}_1} |\langle \nu(P_\epsilon) - \nu(P), h \rangle_{\mathcal{H}} + P_\epsilon\dot{\nu}_P^*(h)| + o(\epsilon). \end{aligned}$$

The second term on the right is $o(\epsilon)$ by (S13), and so it remains to show that the leading term is also $o(\epsilon)$. In fact, we will have shown that the leading term is zero if we can show that ϕ_P is P_ϵ -Bochner integrable, since that would imply that $\langle P_\epsilon\phi_P, h \rangle_{\mathcal{H}} = P_\epsilon\langle \phi_P, h \rangle_{\mathcal{H}} = P_\epsilon\dot{\nu}_P^*(h)$. To see that ϕ_P is indeed P_ϵ -Bochner integrable, note that

$$\begin{aligned} \int \|\phi_P(z)\|_{\mathcal{H}} P_\epsilon(dz) &= \int \|\phi_P(z)\|_{\mathcal{H}} \frac{dP_\epsilon}{dP}(z) P(dz) \\ &\leq \left| \int \|\phi_P(z)\|_{\mathcal{H}} \left[\frac{dP_\epsilon}{dP}(z) - 1 \right] P(dz) \right| + \int \|\phi_P(z)\|_{\mathcal{H}} P(dz) \\ &\leq \|\phi_P\|_{L^2(P; \mathcal{H})} \left\| \frac{dP_\epsilon}{dP} - 1 \right\|_{L^2(P)} + \|\phi_P\|_{L^2(P; \mathcal{H})} \\ &\leq \|\phi_P\|_{L^2(P; \mathcal{H})} \left(\left\| \frac{dP_\epsilon}{dP} - 1 - \epsilon s \right\|_{L^2(P)} + \epsilon \|s\|_{L^2(P)} \right) + \|\phi_P\|_{L^2(P; \mathcal{H})}, \end{aligned}$$

where the first inequality holds by the triangle inequality, the second by Cauchy Schwarz and Jensen's

inequality, and the third by the triangle inequality. The right-hand side is finite since ϕ_P is P -Bochner square integrable and (S12) holds. \square

Proofs for Section 2.4

Proof of Lemma 1. We begin by showing that $\sum_{k=1}^{\infty} \beta_k^2 P \dot{\nu}_P(h_k)^2 < \infty$ and that $(\beta_k \dot{\nu}_P^*(h_k)(z))_{k=1}^{\infty}$ is P -a.s. a square summable sequence. To see why this is the case, note that, by the monotone convergence theorem,

$$\begin{aligned} E_P \left[\sum_{k=1}^{\infty} \beta_k^2 \dot{\nu}_P^*(h_k)(Z)^2 \right] &= \lim_{K \rightarrow \infty} E_P \left[\sum_{k=1}^K \beta_k^2 \dot{\nu}_P^*(h_k)(Z)^2 \right] = \lim_{K \rightarrow \infty} \sum_{k=1}^K \beta_k^2 \|\dot{\nu}_P^*(h_k)\|_{L^2(P)}^2 \\ &\leq \lim_{K \rightarrow \infty} \sum_{k=1}^K \beta_k^2 \|\dot{\nu}_P^*\|_{\text{op}}^2 \|h_k\|_{\mathcal{H}}^2 \leq \|\dot{\nu}_P^*\|_{\text{op}}^2 \lim_{K \rightarrow \infty} \sum_{k=1}^K \beta_k^2 < \infty, \end{aligned} \quad (\text{S14})$$

where the final inequality holds because $\dot{\nu}_P^* : \mathcal{H} \rightarrow L^2(P)$ is a bounded operator and $(\beta_k)_{k=1}^{\infty}$ is square summable. The above implies that $\sum_{k=1}^{\infty} \beta_k^2 \dot{\nu}_P^*(h_k)(z)^2$ is finite on a P -probability one set \mathcal{Z}^{β} . Hence, $\phi_P^{\beta}(z) \in \mathcal{H}$ P -a.s. To see that $\phi_P^{\beta} \in L^2(P; \mathcal{H})$, note that, by the continuity and linearity of inner products and the orthonormality of $(h_k)_{k=1}^{\infty}$,

$$\begin{aligned} &\int \langle \phi_P^{\beta}(z), \phi_P^{\beta}(z) \rangle_{\mathcal{H}} P(dz) \\ &= \int_{\mathcal{Z}^{\beta}} \left\langle \lim_{K \rightarrow \infty} \sum_{k=1}^K \beta_k \dot{\nu}_P^*(h_k)(z) h_k, \lim_{K' \rightarrow \infty} \sum_{k'=1}^{K'} \beta_{k'} \dot{\nu}_P^*(h_{k'})(z) h_{k'} \right\rangle_{\mathcal{H}} P(dz) \\ &= \int_{\mathcal{Z}^{\beta}} \lim_{K \rightarrow \infty} \lim_{K' \rightarrow \infty} \left\langle \sum_{k=1}^K \beta_k \dot{\nu}_P^*(h_k)(z) h_k, \sum_{k'=1}^{K'} \beta_{k'} \dot{\nu}_P^*(h_{k'})(z) h_{k'} \right\rangle_{\mathcal{H}} P(dz) \\ &= \int_{\mathcal{Z}^{\beta}} \lim_{K \rightarrow \infty} \lim_{K' \rightarrow \infty} \sum_{k=1}^K \sum_{k'=1}^{K'} \beta_k \dot{\nu}_P^*(h_k)(z) \beta_{k'} \dot{\nu}_P^*(h_{k'})(z) \langle h_k, h_{k'} \rangle_{\mathcal{H}} P(dz) \\ &= \int_{\mathcal{Z}^{\beta}} \lim_{K \rightarrow \infty} \sum_{k=1}^K \beta_k^2 \dot{\nu}_P^*(h_k)(z)^2 P(dz) = E_P \left[\sum_{k=1}^{\infty} \beta_k^2 \dot{\nu}_P^*(h_k)(Z)^2 \right], \end{aligned}$$

which is finite by (S14). It remains to show that, for all $z \in \mathcal{Z}^{\beta}$ and $h \in \mathcal{H}$, $r_P^{\beta}(h)(z) = \langle \phi_P^{\beta}(z), h \rangle_{\mathcal{H}}$. This can be seen by noting that, for any $z \in \mathcal{Z}^{\beta}$ and $h \in \mathcal{H}$,

$$\begin{aligned} &\langle \phi_P^{\beta}(z), h \rangle_{\mathcal{H}} \\ &= \left\langle \phi_P^{\beta}(z), \sum_{k'=1}^{\infty} \langle h, h_{k'} \rangle_{\mathcal{H}} h_{k'} \right\rangle_{\mathcal{H}} = \left\langle \sum_{k=1}^{\infty} \beta_k \dot{\nu}_P^*(h_k)(z) h_k, \sum_{k'=1}^{\infty} \langle h, h_{k'} \rangle_{\mathcal{H}} h_{k'} \right\rangle_{\mathcal{H}} \\ &= \sum_{k=1}^{\infty} \sum_{k'=1}^{\infty} \langle h, h_{k'} \rangle_{\mathcal{H}} \beta_k \dot{\nu}_P^*(h_k)(z) \langle h_k, h_{k'} \rangle_{\mathcal{H}} = \sum_{k=1}^{\infty} \langle h, h_k \rangle_{\mathcal{H}} \beta_k \dot{\nu}_P^*(h_k)(z) = r_P^{\beta}(h)(z). \end{aligned}$$

As $h \in \mathcal{H}$ and $z \in \mathcal{Z}^{\beta}$ were arbitrary, $r_P^{\beta}(\cdot)(z)$ is a bounded linear functional with Riesz representation $\phi_P^{\beta}(z)$ for all $z \in \mathcal{Z}^{\beta}$. \square

Lemma S7. *Fix a score s in the tangent set of \mathcal{P} at P and let $\{P_{\epsilon} : \epsilon \in [0, \delta)\} \in \mathcal{P}(P, \mathcal{P}, s)$ satisfy (S12). If ν is pathwise differentiable at P , then*

$$\left\| \nu(P) - \nu(P_{\epsilon}) + P_{\epsilon} \phi_P^{\beta} - \sum_{k=1}^{\infty} (1 - \beta_k) \langle \nu(P) - \nu(P_{\epsilon}), h_k \rangle_{\mathcal{H}} h_k \right\|_{\mathcal{H}}$$

$$= \left(1 + \|\phi_P^\beta\|_{L^2(P; \mathcal{H})}\right) \cdot o(\epsilon),$$

where the $o(\epsilon)$ terms do not depend on the choice of β .

Proof of Lemma S7. Suppose that s and $\{P_\epsilon : \epsilon \in [0, \delta)\}$ are as in the statement of the lemma and that ν is pathwise differentiable at P . Let \mathcal{H}_1 denote the unit ball of \mathcal{H} . Since $\nu(P_\epsilon) - \nu(P_0) = \sum_{k=1}^{\infty} \langle \nu(P_\epsilon) - \nu(P), h_k \rangle_{\mathcal{H}} h_k$, it holds that

$$\begin{aligned} \nu(P) - \nu(P_\epsilon) + P_\epsilon \phi_P^\beta - \sum_{k=1}^{\infty} (1 - \beta_k) \langle \nu(P) - \nu(P_\epsilon), h_k \rangle_{\mathcal{H}} h_k \\ = P_\epsilon \phi_P^\beta - \sum_{k=1}^{\infty} \beta_k \langle \nu(P_\epsilon) - \nu(P), h_k \rangle_{\mathcal{H}} h_k. \end{aligned}$$

The remainder of our analysis bounds the terms on the right of the following decomposition, which holds by the triangle inequality:

$$\begin{aligned} & \left\| P_\epsilon \phi_P^\beta - \sum_{k=1}^{\infty} \beta_k \langle \nu(P_\epsilon) - \nu(P), h_k \rangle_{\mathcal{H}} h_k \right\|_{\mathcal{H}} \\ & \leq \left\| P_\epsilon \phi_P^\beta - \epsilon \int s(z) \phi_P^\beta(z) P(dz) \right\|_{\mathcal{H}} \\ & \quad + \epsilon \left\| \int s(z) \phi_P^\beta(z) P(dz) - \sum_{k=1}^{\infty} \beta_k \langle \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} h_k \right\|_{\mathcal{H}} \\ & \quad + \left\| \sum_{k=1}^{\infty} \beta_k \langle \epsilon \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} h_k - \sum_{k=1}^{\infty} \beta_k \langle \nu(P_\epsilon) - \nu(P), h_k \rangle_{\mathcal{H}} h_k \right\|_{\mathcal{H}}. \end{aligned} \tag{S15}$$

To bound the leading term, we use that $\phi_P^\beta \in L^2(P; \mathcal{H})$ by Lemma 1 and that $P \phi_P^\beta = 0$, which give that

$$\begin{aligned} \left\| P_\epsilon \phi_P^\beta - \epsilon \int s(z) \phi_P^\beta(z) P(dz) \right\|_{\mathcal{H}} &= \sup_{h \in \mathcal{H}_1} \left\langle \int \left(\frac{dP_\epsilon}{dP}(z) - 1 - \epsilon s(z) \right) \phi_P^\beta(z) P(dz), h \right\rangle_{\mathcal{H}} \\ &= \sup_{h \in \mathcal{H}_1} \int \left(\frac{dP_\epsilon}{dP}(z) - 1 - \epsilon s(z) \right) \langle \phi_P^\beta(z), h \rangle_{\mathcal{H}} P(dz). \end{aligned}$$

By twice applying Cauchy-Schwarz, once in \mathcal{H} and once in $L^2(P)$, and recalling (S12) and that \mathcal{H}_1 is the unit ball of \mathcal{H} , the right-hand side can be seen to be upper bounded by $\|\phi_P^\beta\|_{L^2(P)} \cdot o(\epsilon)$, where the $o(\epsilon)$ term denotes the behavior of the term on the left-hand side of (S12), which does not depend on β .

We will show that the second term on the right of (S15) is zero. To do this, we recall that (i) $\phi_P^\beta \in L^2(P; \mathcal{H})$ by Lemma 1, (ii) $\phi_P^\beta := \sum_{k=1}^{\infty} \beta_k \dot{\nu}_P^*(h_k)(z) h_k$, (iii) inner products are continuous, (iv) $(h_k)_{k=1}^{\infty}$ is an orthonormal basis of \mathcal{H}_k , and (v) $\dot{\nu}_P^*$ is the adjoint of $\dot{\nu}_P$. Applying these facts in sequence justifies the following for any basis element h_k :

$$\begin{aligned} & \left\langle \int s(z) \phi_P^\beta(z) P(dz), h_k \right\rangle_{\mathcal{H}} \\ &= \int s(z) \langle \phi_P^\beta(z), h_k \rangle_{\mathcal{H}} P(dz) = \int s(z) \left\langle \sum_{k'=1}^{\infty} \beta_{k'} \dot{\nu}_P^*(h_{k'})(z) h_{k'}, h_k \right\rangle_{\mathcal{H}} P(dz) \\ &= \int s(z) \left[\sum_{k'=1}^{\infty} \beta_{k'} \dot{\nu}_P^*(h_{k'})(z) \langle h_{k'}, h_k \rangle_{\mathcal{H}} \right] P(dz) \\ &= \beta_k \langle h_k, h_k \rangle_{\mathcal{H}} \int s(z) \dot{\nu}_P^*(h_k)(z) P(dz) = \beta_k \langle h_k, h_k \rangle_{\mathcal{H}} \langle \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} \end{aligned}$$

$$= \left\langle \sum_{k'=1}^{\infty} \beta_{k'} \langle \dot{\nu}_P(s), h_{k'} \rangle_{\mathcal{H}} h_{k'}, h_k \right\rangle_{\mathcal{H}}$$

As h_k was an arbitrary element of an orthonormal basis of \mathcal{H} , this shows that the second term on the right of (S15) is zero.

We will show that the third term on the right of (S15) is $o(\epsilon)$. We begin by noting that

$$\begin{aligned} & \left\| \sum_{k=1}^{\infty} \beta_k \langle \epsilon \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} h_k - \sum_{k=1}^{\infty} \beta_k \langle \nu(P_\epsilon) - \nu(P), h_k \rangle_{\mathcal{H}} h_k \right\|_{\mathcal{H}} \\ &= \left\| \sum_{k=1}^{\infty} \beta_k \langle \nu(P_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} h_k \right\|_{\mathcal{H}} \\ &= \sup_{h \in \mathcal{H}_1} \left\langle \sum_{k=1}^{\infty} \beta_k \langle \nu(P_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} h_k, h \right\rangle_{\mathcal{H}} \\ &= \sup_{h \in \mathcal{H}_1} \sum_{k=1}^{\infty} \beta_k \langle \nu(P_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} \langle h_k, h \rangle_{\mathcal{H}}. \end{aligned}$$

Combining the above with the fact that all β_k belong to $[0, 1]$, the Cauchy-Schwarz inequality in ℓ^2 , and Parseval's identity, the above shows that

$$\begin{aligned} & \left\| \sum_{k=1}^{\infty} \beta_k \langle \epsilon \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} h_k - \sum_{k=1}^{\infty} \beta_k \langle \nu(P_\epsilon) - \nu(P), h_k \rangle_{\mathcal{H}} h_k \right\|_{\mathcal{H}} \\ &\leq \sup_{h \in \mathcal{H}_1} \sum_{k=1}^{\infty} |\langle \nu(P_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} \langle h_k, h \rangle_{\mathcal{H}}| \\ &\leq \left(\sum_{k=1}^{\infty} \langle \nu(P_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}}^2 \right)^{1/2} \left(\sup_{h \in \mathcal{H}_1} \sum_{k=1}^{\infty} \langle h_k, h \rangle_{\mathcal{H}}^2 \right)^{1/2} \\ &= \|\nu(P_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s)\|_{\mathcal{H}}. \end{aligned}$$

The right-hand side does not depend on β and, by the pathwise differentiability of ν , is $o(\epsilon)$. This completes the proof. \square

C.2 Proofs for Section 3

We now prove the sufficient condition for pathwise differentiability that we presented in the main text. We refer the interested reader to Remark 2 in Appendix A.5 of [Bickel et al. \(1993\)](#) for an alternative characterization of pathwise differentiability that may also be useful in some contexts.

Proof of Lemma 2. Suppose that (i) and (ii) hold. Let $\{P_\epsilon : \epsilon \in [0, \delta]\} \in \mathcal{P}(P, \mathcal{P}, s)$ for some s in the tangent set of \mathcal{P} at P . Since $\mathcal{S}(P)$ is dense in $\dot{\mathcal{P}}_P$, there exists a $\mathcal{S}(P)$ -valued sequence $(s_n)_{n=1}^{\infty}$ such that $s_n \rightarrow s$ in $L^2(P)$. For each n , let $\{P_\epsilon^{[n]} : \epsilon\}$ be the element of $\mathcal{P}(P, \mathcal{P}, s_n)$ satisfying $\|\nu(P_\epsilon^{[n]}) - \nu(P) - \epsilon \eta_P(s)\|_{\mathcal{H}} = o(\epsilon)$ that is guaranteed to exist by (i). Observe that, for any $\epsilon > 0$ and any map $N : (0, \delta] \rightarrow \mathbb{N}$, the triangle inequality and the linearity of η_P show that

$$\begin{aligned} & \|\epsilon^{-1} [\nu(P_\epsilon) - \nu(P)] - \eta_P(s)\|_{\mathcal{H}} \\ &\leq \left\| \epsilon^{-1} \left[\nu(P_\epsilon^{[N(\epsilon)]}) - \nu(P) \right] - \eta_P(s_{N(\epsilon)}) \right\|_{\mathcal{H}} \\ &\quad + \epsilon^{-1} \left\| \nu(P_\epsilon^{[N(\epsilon)]}) - \nu(P_\epsilon) \right\|_{\mathcal{H}} + \|\eta_P(s_{N(\epsilon)} - s)\|_{\mathcal{H}}. \end{aligned} \tag{S16}$$

Our argument will be based on the above with a map $N : (0, \delta] \rightarrow \mathbb{N}$ that we construct to satisfy all of the following properties as $\epsilon \rightarrow 0$:

- a) $N(\epsilon) \rightarrow \infty$;
- b) $\|\epsilon^{-1}[\sqrt{dP_\epsilon^{[N(\epsilon)]}} - \sqrt{dP}] - \frac{1}{2}s_{N(\epsilon)}\sqrt{dP}\|_{L^2(\lambda)} \rightarrow 0$;
- c) $\|\epsilon^{-1}[\nu(P_\epsilon^{[N(\epsilon)]}) - \nu(P)] - \eta_P(s_{N(\epsilon)})\|_{\mathcal{H}} \rightarrow 0$.

We will first establish that such a map exists, and then show that, when combined with (S16), this establishes that ν is pathwise differentiable at P with local parameter $\dot{\nu}_P = \eta_P$.

Now, for any fixed $n \in \mathbb{N}$, the quadratic mean differentiability of $\{P_\epsilon^{[n]} : \epsilon\}$ paired with the fact that $\|\nu(P_\epsilon^{[n]}) - \nu(P) - \epsilon\eta_P(s)\|_{\mathcal{H}} = o(\epsilon)$ implies that $\lim_{\epsilon \rightarrow 0} f_n(\epsilon) = 0$, where

$$\begin{aligned} f_n(\epsilon) := & \left\| \epsilon^{-1} \left[\sqrt{dP_\epsilon^{[n]}} - \sqrt{dP} \right] - \frac{1}{2}s_n\sqrt{dP} \right\|_{L^2(\lambda)} \\ & + \left\| \epsilon^{-1} \left[\nu(P_\epsilon^{[n]}) - \nu(P) \right] - \eta_P(s_n) \right\|_{\mathcal{H}}. \end{aligned}$$

We use this fact to define $N : (0, \delta] \rightarrow \mathbb{N}$ via a recursive formulation. For a strictly decreasing positive sequence $(\epsilon_\ell)_{\ell=1}^\infty$ that we will define momentarily, we let $N(\epsilon) := \max\{\ell \in \mathbb{N} : \epsilon \leq \epsilon_\ell\}$. This sequence will be constructed so that $\epsilon_\ell \downarrow 0$ as $\ell \uparrow \infty$, which ensures that the maximum used to define $N(\epsilon)$ is well-defined and that $N(\epsilon) \rightarrow \infty$ as $\epsilon \rightarrow 0$ or, in other words, condition a) holds. The construction of the sequence will also ensure that $f_{N(\epsilon)}(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$, which will guarantee that conditions b) and c) hold as well. We now construct this sequence. We first let $\epsilon_1 := \delta$. Then, recursively from $\ell = 1, 2, \dots$, we let $\epsilon_{\ell+1} := \frac{1}{2} \sup\{\epsilon \in (0, \epsilon_\ell] : \sup_{\epsilon' \in (0, \epsilon]} f_{\ell+1}(\epsilon') \leq 2^{-(\ell+1)}\}$; since $f_{\ell+1}(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$, $\epsilon_{\ell+1}$ is well-defined and positive. Also, by its definition, $\epsilon_{\ell+1} \leq \epsilon_\ell/2$ and $f_{\ell+1}(\epsilon) \leq 2^{-(\ell+1)}$ for all $\epsilon \leq \epsilon_{\ell+1}$. As a consequence, $f_{N(\epsilon)}(\epsilon) \leq 2^{-N(\epsilon)}$ for all $\epsilon \leq \epsilon_2$. Since $N(\epsilon) \rightarrow \infty$ as $\epsilon \rightarrow 0$, this implies that $f_{N(\epsilon)}(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$, which implies that b) and c) hold.

Having now defined $N : (0, \delta] \rightarrow \mathbb{N}$, we return to (S16). Because s was an arbitrary element of the tangent set, we will have established that ν is pathwise differentiable at P with local parameter $\dot{\nu}_P = \eta_P$ if we can show that the right-hand side of that display converges to zero as $\epsilon \rightarrow 0$. Since the choice of N ensured that c) holds, the first term on the right-hand side of that display goes to zero as $\epsilon \rightarrow 0$. Since $\eta_P : \dot{\mathcal{P}}_P \rightarrow \mathcal{H}$ is a bounded linear operator by (i) and $\lim_{\epsilon \rightarrow 0} s_{N(\epsilon)} = s$ by virtue of the fact that $\lim_{n \rightarrow \infty} s_n = s$ and $N(\epsilon) \rightarrow \infty$ as $\epsilon \rightarrow 0$, the third term on the right-hand side of that display goes to zero as $\epsilon \rightarrow 0$. It remains to study the second term. For this term, we will first show that the Hellinger distance between $P_\epsilon^{[N(\epsilon)]}$ and P_ϵ is $o(\epsilon)$, and then we will leverage the local Lipschitz property of ν that holds by (ii). Beginning by studying the Hellinger distance, we use the triangle inequality to show that

$$\begin{aligned} H(P_\epsilon^{[N(\epsilon)]}, P_\epsilon) &:= \left\| \sqrt{dP_\epsilon^{[N(\epsilon)]}} - \sqrt{dP_\epsilon} \right\|_{L^2(\lambda)} \\ &\leq \left\| \sqrt{dP_\epsilon} - \sqrt{dP} - \frac{1}{2}\epsilon s\sqrt{dP} \right\|_{L^2(\lambda)} \\ &\quad + \left\| \sqrt{dP_\epsilon^{[N(\epsilon)]}} - \sqrt{dP} - \frac{1}{2}\epsilon s\sqrt{dP} \right\|_{L^2(\lambda)} + \frac{1}{2}\epsilon \|s_{N(\epsilon)} - s\|_{L^2(P)}. \end{aligned}$$

The first term on the right is $o(\epsilon)$ by the quadratic mean differentiability of $\{P_\epsilon : \epsilon \in [0, \delta]\}$, the second is $o(\epsilon)$ by b), and the third is $o(\epsilon)$ by the fact that $s_{N(\epsilon)} \rightarrow s$ in $L^2(P)$. Hence, $H(P_\epsilon^{[N(\epsilon)]}, P_\epsilon)$ is $o(\epsilon)$. Letting c be the constant from (ii), this implies that, for all ϵ small enough, the second term in (S16) bounds as follows:

$$\epsilon^{-1} \left\| \nu(P_\epsilon^{[N(\epsilon)]}) - \nu(P_\epsilon) \right\|_{\mathcal{H}} \leq \epsilon^{-1} c H(P_\epsilon^{[N(\epsilon)]}, P_\epsilon) = o(1).$$

□

We now establish a lemma concerning the preservation of quadratic mean differentiability, which we used several times in the derivations for the examples provided in Appendix B.

Lemma S8. *Let \mathcal{P} be a statistical model of distributions of $Z = (X, Y)$ that are equivalent in that, for all $P_1, P_2 \in \mathcal{P}$, $P_1 \ll P_2$ and $P_2 \ll P_1$. Let $\{P_\epsilon : \epsilon \in [0, \delta]\} \in \mathcal{P}(P, \mathcal{P}, s)$, where s is in the tangent set of \mathcal{P} at P . Let $q_\epsilon = \sqrt{\frac{dP_\epsilon}{dP}}$ and $q_{\epsilon, X}, q_{\epsilon, Y|X}$ be the square root of the marginal density and conditional density of P_ϵ relative to P , respectively. Let $s_X(x) = E_P[s(Z) | X = x]$ and $s_{Y|X}(y | x) = s(y) - s_X(x)$. Then, both of the following hold:*

$$\left\| q_{\epsilon, X} - 1 - \frac{\epsilon}{2} s_X \right\|_{L^2(P)} = o(\epsilon), \quad \left\| q_{\epsilon, Y|X} - 1 - \frac{\epsilon}{2} s_{Y|X} \right\|_{L^2(P)} = o(\epsilon).$$

Proof of Lemma S8. Applying Proposition A.5.5 in Bickel et al. (1993) and taking $(x, y) \mapsto x$ as the statistic, we have the first inequality, namely that $\|q_{\epsilon, X} - 1 - \frac{\epsilon}{2} s_X\|_{L^2(P)} = o(\epsilon)$.

We now establish the second equality. Let $s_X^{(\epsilon)} := s_X 1\{|s_X| \leq \epsilon^{-1/2}\}$ and $f(\epsilon) := \|q_{\epsilon, Y|X} - 1 - \frac{\epsilon}{2} s_{Y|X}\|_{L^2(P)}^2$. We have that

$$\begin{aligned} f(\epsilon) &= \int \left(q_{\epsilon, Y|X} [1 - q_{\epsilon, X}] + q_\epsilon - 1 - \frac{\epsilon}{2} s_{Y|X} \right)^2 dP \\ &= \int \left(q_{\epsilon, Y|X} \left[1 - q_{\epsilon, X} + \frac{\epsilon}{2} s_X^{(\epsilon)} \right] + q_\epsilon - 1 - \frac{\epsilon}{2} [s_{Y|X} + s_X^{(\epsilon)}] \right. \\ &\quad \left. - \frac{\epsilon}{2} \left[q_{\epsilon, Y|X} - 1 - \frac{\epsilon}{2} s_{Y|X} \right] s_X^{(\epsilon)} - \frac{\epsilon^2}{4} s_{Y|X} s_X^{(\epsilon)} \right)^2 dP \\ &\leq 4 \int \left(q_{\epsilon, Y|X} \left[1 - q_{\epsilon, X} + \frac{\epsilon}{2} s_X^{(\epsilon)} \right] \right)^2 dP + 4 \int \left(q_\epsilon - 1 - \frac{\epsilon}{2} [s_{Y|X} + s_X^{(\epsilon)}] \right)^2 dP \\ &\quad + \epsilon^2 \int \left[q_{\epsilon, Y|X} - 1 - \frac{\epsilon}{2} s_{Y|X} \right]^2 (s_X^{(\epsilon)})^2 dP + \frac{\epsilon^4}{4} \int s_{Y|X}^2 (s_X^{(\epsilon)})^2 dP, \end{aligned} \tag{S17}$$

where the final inequality uses that $(a + b + c + d)^2 \leq 4(a^2 + b^2 + c^2 + d^2)$. We consider each of the four terms above separately, showing that the first two are $o(\epsilon^2)$, the third is no more than $\epsilon f(\epsilon)$ and the last is $O(\epsilon^3)$. Subtracting the third term from both sides and dividing both sides by $1 - \epsilon$ will then show that $f(\epsilon) = o(\epsilon^2)$. For the first term, note that

$$\begin{aligned} &4 \int \left(q_{\epsilon, Y|X} \left[1 - q_{\epsilon, X} + \frac{\epsilon}{2} s_X^{(\epsilon)} \right] \right)^2 dP \\ &= 4 \iint \left(1 - q_{\epsilon, X}(x) + \frac{\epsilon}{2} s_X^{(\epsilon)}(x) \right)^2 P_{\epsilon, Y|X}(dy | x) P_X(dx) \\ &= 4 \int \left(1 - q_{\epsilon, X} + \frac{\epsilon}{2} s_X^{(\epsilon)} \right)^2 dP_X \\ &\leq 8 \int \left(1 - q_{\epsilon, X} + \frac{\epsilon}{2} s_X \right)^2 dP_X + 8\epsilon^2 \int s_X^2 1\{|s_X| > \epsilon^{-1/2}\} dP, \end{aligned}$$

where we used that $(a + b)^2 \leq 2(a^2 + b^2)$. The first term on the right is $o(\epsilon^2)$ since quadratic mean differentiability is preserved under marginalization as proved earlier, and the second is $o(\epsilon^2)$ by the dominated convergence theorem. For the second term in (S17), similar arguments show that

$$4 \int \left(q_\epsilon - 1 - \frac{\epsilon}{2} [s_{Y|X} + s_X^{(\epsilon)}] \right)^2 dP = o(\epsilon^2).$$

For the third integral in (S17),

$$\epsilon^2 \int \left[q_{\epsilon, Y|X} - 1 - \frac{\epsilon}{2} s_{Y|X} \right]^2 (s_X^{(\epsilon)})^2 dP \leq \epsilon \int \left[q_{\epsilon, Y|X} - 1 - \frac{\epsilon}{2} s_{Y|X} \right]^2 dP = \epsilon f(\epsilon)$$

and, for the final term in (S17),

$$\frac{\epsilon^4}{4} \int s_{Y|X}^2 (s_X^{(\epsilon)})^2 dP \leq \frac{\epsilon^3}{4} \int s_{Y|X}^2 dP.$$

This completes the proof. \square

C.3 Proofs for Section 4

C.3.1 Proofs for Section 4.1

Proof of Lemma 3. Fix $\delta > 0$. Suppose that $\|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$. We will show that

$$\lim_{n \rightarrow \infty} P_0^n \{ \|\mathcal{D}_n^1\|_{\mathcal{H}} > n^{-1/2} \delta \} = 0. \quad (\text{S18})$$

As δ was arbitrary, this will show that $\mathcal{D}_n^1 = o_p(n^{-1/2})$. An analogous argument can be used to show that $\|\phi_n^2 - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$ implies that $\mathcal{D}_n^2 = o_p(n^{-1/2})$.

Let $1_{\mathcal{E}_n}$ denote the indicator that the event \mathcal{E}_n that $\|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})}^2 \leq \delta^2/2$ and let \mathcal{E}_n^c denote the complement of \mathcal{E}_n . We will leverage the following decomposition when showing (S18):

$$\begin{aligned} P_0^n \{ \|\mathcal{D}_n^1\|_{\mathcal{H}} > n^{-1/2} \delta \} &\leq P_0^n \left(\{ \|\mathcal{D}_n^1\|_{\mathcal{H}} > n^{-1/2} \delta \} \cap \mathcal{E}_n \right) + P_0^n(\mathcal{E}_n^c) \\ &= E_0^n \left[1_{\mathcal{E}_n} P_0^n \left\{ \|\mathcal{D}_n^1\|_{\mathcal{H}} > n^{-1/2} \delta \mid Z_1, \dots, Z_{n/2} \right\} \right] + o(1), \end{aligned} \quad (\text{S19})$$

where E_0^n denotes an expectation under sampling from the n -fold product measure P_0^n and $o(1)$ denotes a deterministic term that goes to zero as $n \rightarrow \infty$. The equality above holds by the law of total probability, the fact that \mathcal{E}_n is measurable with respect to the σ -field generated by $Z_1, \dots, Z_{n/2}$, and the assumption that $\|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$ implies that $P_0^n(\mathcal{E}_n^c) = o(1)$. To show (S18), the above shows that it suffices to show that the first term on the right-hand side is $o(1)$. To this end, note that Chebyshev's inequality for Hilbert-valued random variables (Grenander, 1963) and the bilinearity of inner products shows that

$$\begin{aligned} &1_{\mathcal{E}_n} P_0^n \left\{ \|\mathcal{D}_n^1\|_{\mathcal{H}} > n^{-1/2} \delta \mid Z_1, \dots, Z_{n/2} \right\} \\ &\leq 1_{\mathcal{E}_n} \frac{E_{P_0^n} [\|(P_n^1 - P_0)(\phi_n^1 - \phi_0)\|_{\mathcal{H}}^2 \mid Z_1, \dots, Z_{n/2}]}{n^{-1} \delta^2} \\ &= 1_{\mathcal{E}_n} \frac{(n/2)^{-1} E_{P_0^n} [P_n^1 \|(I - P_0)(\phi_n^1 - \phi_0)\|_{\mathcal{H}}^2 \mid Z_1, \dots, Z_{n/2}]}{n^{-1} \delta^2} \\ &\quad + 1_{\mathcal{E}_n} \frac{4}{n^2} \sum_{(i,j) \in \{n/2+1, \dots, n\}^2 : i \neq j} \frac{E_{P_0^n} [\langle (I - P_0)(\phi_n^1 - \phi_0)(Z_i), (I - P_0)(\phi_n^1 - \phi_0)(Z_k) \rangle_{\mathcal{H}} \mid Z_1, \dots, Z_{n/2}]}{n^{-1} \delta^2}, \end{aligned} \quad (\text{S20})$$

where $(I - P_0)(\phi_n^1 - \phi_0)$ denotes the map $z \mapsto (\phi_n^1 - \phi_0)(z) - P_0(\phi_n^1 - \phi_0)$ and each expectation in the summand on the right-hand side above is well-defined since $(\phi_n^1 - \phi_0) \in L^2(P_0; \mathcal{H})$ implies that $(I - P_0)(\phi_n^1 - \phi_0) \in L^2(P_0; \mathcal{H})$ as well. In fact, each expectation in the summand on the right-hand side is zero since, by the fact that $(I - P_0)(\phi_n^1 - \phi_0) \in L^2(P_0; \mathcal{H})$ and Fubini's theorem,

$$E_{P_0^n} [\langle (I - P_0)(\phi_n^1 - \phi_0)(Z_i), (I - P_0)(\phi_n^1 - \phi_0)(Z_k) \rangle_{\mathcal{H}} \mid Z_1, \dots, Z_{n/2}]$$

$$\begin{aligned}
&= \int \langle (I - P_0)(\phi_n^1 - \phi_0)(z_1), (I - P_0)(\phi_n^1 - \phi_0)(z_2) \rangle_{\mathcal{H}} P_0^2(dz_1, dz_2) \\
&= \iint \langle (I - P_0)(\phi_n^1 - \phi_0)(z_1), (I - P_0)(\phi_n^1 - \phi_0)(z_2) \rangle_{\mathcal{H}} P_0(dz_1) P_0(dz_2) \\
&= \int \left\langle \int (I - P_0)(\phi_n^1 - \phi_0)(z_1) P_0(dz_1), (I - P_0)(\phi_n^1 - \phi_0)(z_2) \right\rangle_{\mathcal{H}} P_0(dz_2) \\
&= \int \langle 0, (I - P_0)(\phi_n^1 - \phi_0)(z_2) \rangle_{\mathcal{H}} P_0(dz_2) = 0.
\end{aligned}$$

Returning to (S20) and simplifying the first term on the right-hand side of that expression, this shows that

$$1_{\mathcal{E}_n} P_0^n \left\{ \|\mathcal{D}_n^1\|_{\mathcal{H}} > n^{-1/2} \delta \mid Z_1, \dots, Z_{n/2} \right\} \leq 1_{\mathcal{E}_n} \frac{2\|(I - P_0)(\phi_n^1 - \phi_0)\|_{L^2(P_0; \mathcal{H})}^2}{\delta^2}.$$

Using that $P_0(\phi_n^1 - \phi_0)$ is a minimizer over $h \in \mathcal{H}$ of $\|\phi_n^1 - \phi_0 - h\|_{L^2(P; \mathcal{H})}^2$ and subsequently leveraging the definition of the event \mathcal{E}_n , this shows that

$$\begin{aligned}
&1_{\mathcal{E}_n} P_0^n \left\{ \|\mathcal{D}_n^1\|_{\mathcal{H}} > n^{-1/2} \delta \mid Z_1, \dots, Z_{n/2} \right\} \\
&\leq 1_{\mathcal{E}_n} \frac{2\|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})}^2}{\delta^2} \leq \min \left\{ 1, \frac{2\|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})}^2}{\delta^2} \right\}.
\end{aligned}$$

Note that the right-hand side above is no larger than 1. Taking an expectation of both sides over $Z_1, \dots, Z_{n/2} \stackrel{\text{iid}}{\sim} P_0$ and recalling that $\|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$, the dominated convergence theorem shows that the first term on the right-hand side of (S19) is $o(1)$. This completes the proof. \square

Proof of Theorem 2. In this argument, we will let $\tilde{\mathcal{H}}$ denote the Hilbert space of elements $(h, r) \in \mathcal{H} \times \mathbb{R}$ that is equipped with inner product $\langle (h_1, r_1), (h_2, r_2) \rangle_{\tilde{\mathcal{H}}} = \langle h_1, h_2 \rangle_{\mathcal{H}} + r_1 r_2$. Fix $s \in \dot{\mathcal{P}}_P$. Since $\mathcal{R}_n^j = o_P(n^{-1/2})$ and $\mathcal{D}_n^j = o_P(n^{-1/2})$ for $j \in \{1, 2\}$, (21) shows that

$$\begin{pmatrix} \bar{\nu}_n - \nu(P_0) \\ P_n s \end{pmatrix} = \begin{pmatrix} P_n \phi_0 \\ P_n s \end{pmatrix} + o_p(n^{-1/2}) = P_n \begin{pmatrix} \phi_0 \\ s \end{pmatrix} + o_p(n^{-1/2}). \quad (\text{S21})$$

Moreover, $\|(\phi_0, s)\|_{L^2(P; \tilde{\mathcal{H}})}^2 = \|\phi_0\|_{L^2(P; \mathcal{H})}^2 + \|s\|_{L^2(P)}^2 < \infty$. Hence, by Slutsky's lemma and a central limit theorem for Hilbert-valued random variables (see Examples 1.4.7 and 1.8.5 in [van der Vaart and Wellner, 1996](#)), it holds that

$$n^{1/2} \begin{pmatrix} \bar{\nu}_n - \nu(P_0) \\ P_n s \end{pmatrix} \rightsquigarrow \begin{pmatrix} \mathbb{H} \\ S \end{pmatrix}, \quad (\text{S22})$$

where (\mathbb{H}, S) is a tight $\tilde{\mathcal{H}}$ -valued Gaussian random variable that is such that

$$\langle (h, r), (\mathbb{H}, S) \rangle_{\tilde{\mathcal{H}}} \sim N \left(0, E_0 \left[\{ \langle \phi_0(Z), h \rangle_{\mathcal{H}} + r s(Z) \}^2 \right] \right).$$

Marginalizing the $P_n s$ term on the left-hand side of (S21) shows that (S36) holds.

We will use (S22) along with Theorem 3 in Chapter 5.2 of [Bickel et al. \(1993\)](#) to establish the regularity of $\bar{\nu}_n$. To use this result, it suffices to show that $E[S\mathbb{H}] = \nu_P(s)$. We will establish this by showing that $\langle h, E[S\mathbb{H}] \rangle_{\mathcal{H}} = \langle h, \nu_P(s) \rangle_{\mathcal{H}}$ for all $h \in \mathcal{H}$. To see that this holds, first note that $\langle h, E[S\mathbb{H}] \rangle_{\mathcal{H}} = E[S \langle h, \mathbb{H} \rangle_{\mathcal{H}}]$ since

$$E[\|S\mathbb{H}\|_{\mathcal{H}}^2] = E[S^2 \|\mathbb{H}\|_{\mathcal{H}}^2] \leq E[S^4]^{1/2} E[\|\mathbb{H}\|_{\mathcal{H}}^4]^{1/2} < \infty,$$

where the first inequality holds by Cauchy-Schwarz and the second by Fernique's theorem (Fernique, 1970). Hence, it suffices to show that $E[S\langle h, \mathbb{H} \rangle_{\mathcal{H}}] = \langle h, \dot{\nu}_P(s) \rangle_{\mathcal{H}}$. To show that this is the case, we let $f_h : \tilde{\mathcal{H}} \rightarrow \mathbb{R}^2$ be defined so that $f_h(h_1, r_1) = (\langle h_1, h \rangle_{\mathcal{H}}, r_1)$. For any $(a, b) \in \mathbb{R}^2$, the dot product $(a, b) \cdot f_h(\mathbb{H}, S)$ is equal to $a\langle h, \mathbb{H} \rangle_{\mathcal{H}} + bS = \langle (ah, b), (\mathbb{H}, S) \rangle_{\tilde{\mathcal{H}}}$, which follows a mean-zero normal distribution with variance

$$E_0 \left[\{a\langle \phi_0(Z), h \rangle_{\mathcal{H}} + bS(Z)\}^2 \right] = (a, b)^\top \Sigma (a, b),$$

where

$$\Sigma := \begin{pmatrix} E_0 [\langle \phi_0(Z), h \rangle_{\mathcal{H}}^2] & E_0 [s(Z)\langle \phi_0(Z), h \rangle_{\mathcal{H}}] \\ E_0 [s(Z)\langle \phi_0(Z), h \rangle_{\mathcal{H}}] & E_0 [s(Z)^2] \end{pmatrix}.$$

As $(a, b) \in \mathbb{R}^2$ was arbitrary, it follows that $(\langle h, \mathbb{H} \rangle_{\mathcal{H}}, S) \sim N((0, 0), \Sigma)$. Hence, $E[S\langle h, \mathbb{H} \rangle_{\mathcal{H}}]$ is equal to $E_0 [s(Z)\langle \phi_0(Z), h \rangle_{\mathcal{H}}]$. Finally, note that $\langle \phi_0(Z), h \rangle_{\mathcal{H}} = \dot{\nu}_P^*(h)(Z)$ P -a.s. since ϕ_0 is the EIF of ν , and so

$$E_0 [s(Z)\langle \phi_0(Z), h \rangle_{\mathcal{H}}] = E_0 [s(Z)\dot{\nu}_P^*(h)(Z)] = \langle s, \dot{\nu}_P^*(h) \rangle_{L^2(P)} = \langle h, \dot{\nu}_P(s) \rangle_{\mathcal{H}}.$$

□

Proofs for Section 4.2

Proof of Theorem 3. Suppose the conditions of Theorem 2 hold and that $\|\phi_0\|_{L^2(P_0; \mathcal{H})}$ is strictly positive, $\Omega_n \in \mathcal{O}$, $\Omega_0 \in \mathcal{O}$, and $\|\Omega_n - \Omega_0\|_{\text{op}} = o_p(1)$. For brevity, we will let $w_n(\cdot) := w(\cdot; \Omega_n)$ and $w_0(\cdot) := w(\cdot; \Omega_0)$ in this proof. By Theorem 2, $n^{1/2}[\bar{\nu}_n - \nu(P_0)] \rightsquigarrow \mathbb{H}$, where \mathbb{H} is as defined in that theorem. Slutsky's lemma and the continuous mapping theorem can further be used to show that $n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \rightsquigarrow w_0(\mathbb{H})$. To see this, first note that

$$n \cdot w_n[\bar{\nu}_n - \nu(P_0)] = n \cdot w_0[\bar{\nu}_n - \nu(P_0)] + n \cdot w[\bar{\nu}_n - \nu(P_0); \Omega_n - \Omega_0]. \quad (\text{S23})$$

The first of the two terms on the right converges weakly to $w_0(\mathbb{H})$ by the continuous mapping theorem, where we have used that, by virtue of belonging to \mathcal{O} , Ω_0 is a continuous operator, and therefore $w_0 : \mathcal{H} \rightarrow \mathbb{R}$ is a continuous functional. The second term on the right is $o_p(1)$, since, by Cauchy-Schwarz, the definition of the operator norm, and the continuous mapping theorem,

$$\begin{aligned} |n \cdot w[\bar{\nu}_n - \nu(P_0); \Omega_n - \Omega_0]| &\leq n \|(\Omega_n - \Omega_0)[\bar{\nu}_n - \nu(P_0)]\|_{\mathcal{H}} \|\bar{\nu}_n - \nu(P_0)\|_{\mathcal{H}} \\ &\leq \|\Omega_n - \Omega_0\|_{\text{op}} \|n^{1/2}[\bar{\nu}_n - \nu(P_0)]\|_{\mathcal{H}}^2 = o_p(1)O_p(1) = o_p(1). \end{aligned}$$

Plugging this into (S23) and applying Slutsky's lemma shows that $n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \rightsquigarrow w_0(\mathbb{H})$.

We apply Corollary 3.3 of Bogachev (1996) to show that $w_0(\mathbb{H})$ is absolutely continuous. To apply this corollary, it suffices to show that $w_0 : \mathcal{H} \rightarrow \mathbb{R}$ is locally Lipschitz and that the image of the Gateaux derivative $dw_0(h; \cdot)$ is \mathbb{P} -a.s. equal to \mathbb{R} , where \mathbb{P} is the distribution of \mathbb{H} . For the Gateaux derivative condition, we note that, for any $g \in \mathcal{H}$, $dw_0(h; g) = \frac{d}{de}w(h + eg; \Omega_0) = \langle \Omega_0(g), h \rangle_{\mathcal{H}} + \langle \Omega_0(h), g \rangle_{\mathcal{H}} = 2\langle \Omega_0(g), h \rangle_{\mathcal{H}}$, where the latter equality used that Ω_0 is self-adjoint. Since Ω_0 is positive definite, $\langle \Omega_0(h), h \rangle_{\mathcal{H}} > 0$ for all $h \in \mathcal{H}$. Hence, for any $h \in \mathcal{H} \setminus \{0\}$, the image of $dw_0(h; \cdot)$ is equal to \mathbb{R} ; this can be seen by considering $dw_0(h; ch)$ with c varying over \mathbb{R} . As $\|\phi_0\|_{L^2(P_0; \mathcal{H})} > 0$, $\mathcal{H} \setminus \{0\}$ is a \mathbb{P} -probability one set. The locally Lipschitz property follows from the fact that, for all $h \in \mathcal{H}$ and all g_1

and g_2 in the unit ball \mathcal{H}_1 of \mathcal{H} ,

$$\begin{aligned}
& |w_0(h + g_1) - w_0(h + g_2)| \\
&= \left| \int_0^1 [dw_0(h; \epsilon g_1) - dw_0(h; \epsilon g_2)] d\epsilon \right| \leq \int_0^1 |dw_0(h; \epsilon g_1) - dw_0(h; \epsilon g_2)| d\epsilon \\
&= 2 |\langle \Omega_0(g_1 - g_2), h \rangle_{\mathcal{H}}| \int_0^1 \epsilon d\epsilon = |\langle \Omega_0(g_1 - g_2), h \rangle_{\mathcal{H}}| \leq \|\Omega_0(g_1 - g_2)\|_{\mathcal{H}} \|h\|_{\mathcal{H}} \\
&\leq \|\Omega_0\|_{\text{op}} \|g_1 - g_2\|_{\mathcal{H}} \|h\|_{\mathcal{H}} = \|\Omega_0\|_{\text{op}} \|(h + g_1) - (h + g_2)\|_{\mathcal{H}} \|h\|_{\mathcal{H}}.
\end{aligned}$$

Hence, w_0 is $\|\Omega_0\|_{\text{op}} \|h\|_{\mathcal{H}}$ -Lipschitz continuous in the radius-one ball centered at h . As h was arbitrary, w_0 is locally Lipschitz. Corollary 3.3 of [Bogachev \(1996\)](#) thus shows that $w_0(\mathbb{H})$ is an absolutely continuous random variable.

Because convergence in distribution implies convergence of cumulative distribution functions at continuity points, $n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \rightsquigarrow w_0(\mathbb{H})$ implies that

$$P_0^n \{n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \leq \zeta_{1-\alpha} + \delta\} \xrightarrow{n \rightarrow \infty} \mathbb{F}(\zeta_{1-\alpha} + \delta) \quad (\text{S24})$$

for all $\delta \in \mathbb{R}$, where \mathbb{F} is the cumulative distribution function of $w_0(\mathbb{H})$. In what follows, we will use this fact twice when establishing the asymptotic validity of the $(1 - \alpha)$ -confidence sets $\mathcal{C}_n(\hat{\zeta}_n)$. When doing so, we will also use that the event $\{\nu(P_0) \in \mathcal{C}_n(\hat{\zeta}_n)\}$ is the same as the event $\{n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \leq \hat{\zeta}_n\}$.

We now establish (ii). To do this, we use that, for any $\delta > 0$,

$$\begin{aligned}
& P_0^n \left\{ n \cdot w_n[\bar{\nu}_n - \nu(P_0)] > \hat{\zeta}_n \right\} \\
&\leq P_0^n \left\{ n \cdot w_n[\bar{\nu}_n - \nu(P_0)] > \hat{\zeta}_n, \hat{\zeta}_n - \zeta_{1-\alpha} \geq -\delta \right\} + P_0^n \left\{ \hat{\zeta}_n - \zeta_{1-\alpha} < -\delta \right\} \\
&\leq P_0^n \left\{ n \cdot w_n[\bar{\nu}_n - \nu(P_0)] > \zeta_{1-\alpha} - \delta \right\} + P_0^n \left\{ \hat{\zeta}_n - \zeta_{1-\alpha} < -\delta \right\}.
\end{aligned}$$

Subtracting both sides from 1 yields that

$$\begin{aligned}
& P_0^n \left\{ n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \leq \hat{\zeta}_n \right\} \\
&\geq P_0^n \left\{ n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \leq \zeta_{1-\alpha} - \delta \right\} - P_0^n \left\{ \hat{\zeta}_n - \zeta_{1-\alpha} < -\delta \right\}.
\end{aligned}$$

Taking $n \rightarrow \infty$, applying (S24), and then taking $\delta \downarrow 0$ shows that, if $\hat{\zeta}_n$ is an asymptotically conservative estimator of $\zeta_{1-\alpha}$ in the sense stated in (ii), then

$$\liminf_n P_0^n \left\{ n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \leq \hat{\zeta}_n \right\} \geq \mathbb{F}(\zeta_{1-\alpha}) = 1 - \alpha.$$

This establishes (ii).

We now establish (i). To do this, we use that, for any $\delta > 0$,

$$\begin{aligned}
& P_0^n \left\{ n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \leq \hat{\zeta}_n \right\} \\
&\leq P_0^n \left\{ n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \leq \hat{\zeta}_n, \hat{\zeta}_n - \zeta_{1-\alpha} \leq \delta \right\} + P_0^n \left\{ \hat{\zeta}_n - \zeta_{1-\alpha} > \delta \right\} \\
&\leq P_0^n \left\{ n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \leq \zeta_{1-\alpha} + \delta \right\} + P_0^n \left\{ \hat{\zeta}_n - \zeta_{1-\alpha} > \delta \right\}.
\end{aligned}$$

Taking $n \rightarrow \infty$, applying (S24), and then taking $\delta \downarrow 0$ shows that, if $\hat{\zeta}_n$ is a consistent estimator of $\zeta_{1-\alpha}$, then $\limsup_n P_0^n \{n \cdot w_n[\bar{\nu}_n - \nu(P_0)] \leq \hat{\zeta}_n\} \leq \mathbb{F}(\zeta_{1-\alpha}) = 1 - \alpha$. Combining this with (ii) gives (i). \square

In the following result, we write $o_p(1)$ to denote a term that converges to zero in probability marginally over the randomness both in the original sample (Z_1, \dots, Z_n) and the bootstrap sample $(Z_1^\#, \dots, Z_n^\#)$.

Lemma S9. *If $\|\phi_n^j - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$ for each $j \in \{1, 2\}$, then $\|\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#\|_{\mathcal{H}} = o_p(1)$, where $\mathbb{H}_{n,0}^\# := n^{1/2} \frac{1}{2} \sum_{j=1}^2 (P_n^{j,\#} - P_n^j) \phi_0$.*

Proof of Lemma S9. This proof bears resemblance to that of Lemma 3. In what follows we use Z_1^n as shorthand for the sample (Z_1, Z_2, \dots, Z_n) . Note that, for any $\delta > 0$,

$$\begin{aligned} \Pr \left\{ \left\| \mathbb{H}_n^\# - \mathbb{H}_{n,0}^\# \right\|_{\mathcal{H}} > \delta \mid Z_1^n \right\} &= \Pr \left\{ \left\| \frac{1}{2} \sum_{j=1}^2 (P_n^{j,\#} - P_n^j) [\phi_n^j - \phi_0] \right\|_{\mathcal{H}} > n^{-1/2} \delta \mid Z_1^n \right\} \\ &\leq \sum_{j=1}^2 \Pr \left\{ \left\| (P_n^{j,\#} - P_n^j) [\phi_n^j - \phi_0] \right\|_{\mathcal{H}} > n^{-1/2} \delta \mid Z_1^n \right\}. \end{aligned}$$

Taking an expectation of both sides over $Z_1, Z_2, \dots, Z_n \stackrel{\text{iid}}{\sim} P_0$,

$$\Pr \left\{ \left\| \mathbb{H}_n^\# - \mathbb{H}_{n,0}^\# \right\|_{\mathcal{H}} > \delta \right\} = \sum_{j=1}^2 \Pr \left\{ \left\| (P_n^{j,\#} - P_n^j) [\phi_n^j - \phi_0] \right\|_{\mathcal{H}} > n^{-1/2} \delta \right\}.$$

In what follows we show that $\Pr \left\{ \left\| (P_n^{j,\#} - P_n^j) [\phi_n^j - \phi_0] \right\|_{\mathcal{H}} > n^{-1/2} \delta \right\} = o(1)$ when $j = 1$. An analogous argument holds for $j = 2$. As $\delta > 0$ was arbitrary, this will complete the proof.

We begin by noting that

$$\begin{aligned} &\Pr \left\{ \left\| (P_n^{1,\#} - P_n^1) [\phi_n^1 - \phi_0] \right\|_{\mathcal{H}} > n^{-1/2} \delta \mid Z_1^n \right\} \\ &\leq \frac{n}{\delta^2} E \left[\left\| (P_n^{1,\#} - P_n^1) (\phi_n^1 - \phi_0) \right\|_{\mathcal{H}}^2 \mid Z_1^n \right] \\ &= \frac{2}{\delta^2} E \left[P_n^{1,\#} \left\| (I - P_n^1) (\phi_n^1 - \phi_0) \right\|_{\mathcal{H}}^2 \mid Z_1^n \right] \\ &\quad + 4 \sum_{(i,j) \in \{n/2+1, \dots, n\}^2: i \neq j} \frac{E[(I - P_n^1)(\phi_n^1 - \phi_0)(Z_i^\#), (I - P_n^1)(\phi_n^1 - \phi_0)(Z_k^\#)]_{\mathcal{H}} \mid Z_1^n}{n \delta^2}, \end{aligned}$$

where $(I - P_n^1)(\phi_n^1 - \phi_0)$ denotes the map $z \mapsto (\phi_n^1 - \phi_0)(z) - P_n^1(\phi_n^1 - \phi_0)$. Because $Z_i^\#$ and $Z_k^\#$ are independent draws from P_n^1 conditional on Z_1^n when $i \neq k$, each term in the summand on the right-hand side is exactly equal to zero. Since $P_n^{1,\#}$ is the empirical distribution of an iid sample from P_n^1 , the first term on the right rewrites as $(2/\delta^2) \|(I - P_n^1)(\phi_n^1 - \phi_0)\|_{L^2(P_n^1; \mathcal{H})}^2$. As $P_n^1(\phi_n^1 - \phi_0)$ is a minimizer over $h \in \mathcal{H}$ of $\|\phi_n^1 - \phi_0 - h\|_{L^2(P_n^1; \mathcal{H})}^2$, this term upper bounds by $(2/\delta^2) \|\phi_n^1 - \phi_0\|_{L^2(P_n^1; \mathcal{H})}^2$. Plugging this bound into the above yields that

$$\Pr \left\{ \left\| (P_n^{1,\#} - P_n^1) [\phi_n^1 - \phi_0] \right\|_{\mathcal{H}} > n^{-1/2} \delta \mid Z_1^n \right\} \leq \frac{2}{\delta^2} \|\phi_n^1 - \phi_0\|_{L^2(P_n^1; \mathcal{H})}^2.$$

Taking the mean of both sides over the sample $Z_{n/2+1}, Z_{n/2+2}, \dots, Z_n \stackrel{\text{iid}}{\sim} P_0$ used to define P_n^1 shows that

$$\Pr \left\{ \left\| (P_n^{1,\#} - P_n^1) [\phi_n^1 - \phi_0] \right\|_{\mathcal{H}} > n^{-1/2} \delta \mid Z_1, \dots, Z_{n/2} \right\} \leq \frac{2}{\delta^2} \|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})}^2.$$

Using the trivial bound that probabilities are no more than 1 and subsequently taking an expectation on both sides over $Z_1, Z_2, \dots, Z_{n/2} \stackrel{\text{iid}}{\sim} P_0$ yields that

$$\Pr \left\{ \left\| (P_n^{1,\#} - P_n^1) [\phi_n^1 - \phi_0] \right\|_{\mathcal{H}} > n^{-1/2} \delta \right\} \leq E \left[\min \left\{ 1, \frac{2}{\delta^2} \|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})}^2 \right\} \right].$$

Using that $\|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$ by assumption and applying the dominated convergence theorem shows that the right-hand side is $o(1)$, which gives the result. \square

Proof of Theorem 4. By Lemma S9, $\|\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#\|_{\mathcal{H}} = o_p(1)$. By Remark 2.5 of [Giné and Zinn \(1990\)](#) and the fact that $\phi_0 \in L^2(P_0; \mathcal{H})$, $\mathbb{H}_{n,0}^\# \rightsquigarrow \mathbb{H}$ weakly a.s., that is, weakly conditionally on the iid sequence $(Z_i)_{i=1}^\infty$ with P_0 -probability one. By the continuous mapping theorem, this implies that $w(\mathbb{H}_{n,0}^\#; \Omega_0) \rightsquigarrow w(\mathbb{H}; \Omega_0)$ weakly a.s. as well. In what follows we will use these facts, along with the fact that $\|\Omega_n - \Omega_0\|_{\text{op}} = o_p(1)$, to show that $w(\mathbb{H}_n^\#; \Omega_n) \rightsquigarrow w(\mathbb{H}; \Omega_0)$ weakly, conditionally on $(Z_i)_{i=1}^\infty$, in probability, in the sense defined in Chapter 23.2.1 of [van der Vaart \(2000\)](#). To show this, we begin by noting that, as Ω_n and Ω_0 belong to \mathcal{O} ,

$$w(\mathbb{H}_n^\#; \Omega_n) - w(\mathbb{H}_{n,0}^\#; \Omega_0) = w(\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#; \Omega_0) + w(\mathbb{H}_n^\#; \Omega_n - \Omega_0). \quad (\text{S25})$$

We now show that each of the terms on the right are marginally $o_p(1)$. For the first, this follows from the fact that $|w(\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#; \Omega_0)| = |\langle \Omega_0(\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#), \mathbb{H}_n^\# - \mathbb{H}_{n,0}^\# \rangle_{\mathcal{H}}| \leq \|\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#\|_{\mathcal{H}}^2 \|\Omega_0\|_{\text{op}}$, and this upper bound is $o_p(1)$ by Lemma S9. For the second, this follows from the fact that $|w(\mathbb{H}_n^\#; \Omega_n - \Omega_0)| \leq \|\mathbb{H}_n^\#\|_{\mathcal{H}}^2 \|\Omega_n - \Omega_0\|_{\text{op}}$, combined with the fact that $\|\Omega_n - \Omega_0\|_{\text{op}} = o_p(1)$, by assumption, and $\|\mathbb{H}_n^\#\|_{\mathcal{H}}^2 = O_p(1)$, by virtue of the fact that $\|\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#\|_{\mathcal{H}} = o_p(1)$ and $\mathbb{H}_{n,0}^\# \rightsquigarrow \mathbb{H}$ weakly almost surely.

We now derive a form of Slutsky's lemma to show that $w(\mathbb{H}_n^\#; \Omega_0) \rightsquigarrow w(\mathbb{H}; \Omega_0)$ weakly, conditionally on $(Z_i)_{i=1}^\infty$, in probability. In particular, taking $f : \mathbb{R} \rightarrow [-1, 1]$ to be a bounded, 1-Lipschitz function, letting $Z_1^n := (Z_i)_{i=1}^n$, and recalling (S25), we see that

$$\begin{aligned} & |E[f(w(\mathbb{H}_n^\#; \Omega_n)) | Z_1^n] - E[f(w(\mathbb{H}; \Omega_0))]| \\ & \leq |E[f(w(\mathbb{H}_n^\#; \Omega_n)) | Z_1^n] - E[f(w(\mathbb{H}_{n,0}^\#; \Omega_0)) | Z_1^n]| \\ & \quad + |E[f(w(\mathbb{H}_{n,0}^\#; \Omega_0)) | Z_1^n] - E[f(w(\mathbb{H}; \Omega_0))]| \\ & \leq E\left[\min\left\{2, |w(\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#; \Omega_0) + w(\mathbb{H}_n^\#; \Omega_n - \Omega_0)|\right\}\right] |Z_1^n| \\ & \quad + |E[f(w(\mathbb{H}_{n,0}^\#; \Omega_0)) | Z_1^n] - E[f(w(\mathbb{H}; \Omega_0))]|. \end{aligned}$$

Taking a supremum over all 1-Lipschitz $f : \mathbb{R} \rightarrow [-1, 1]$ on both sides and using that $w(\mathbb{H}_{n,0}^\#; \Omega_0) \rightsquigarrow w(\mathbb{H}; \Omega_0)$ weakly a.s. is equivalent to the supremum over such f of the latter term on the right being P_0 -a.s. $o(1)$ (Chapter 23.2.1 of [van der Vaart, 2000](#)), we see that it is P_0 -a.s. true that

$$\begin{aligned} & \sup_f |E[f(w(\mathbb{H}_n^\#; \Omega_n)) | Z_1^n] - E[f(w(\mathbb{H}; \Omega_0))]| \\ & \leq E\left[\min\left\{2, |w(\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#; \Omega_0) + w(\mathbb{H}_n^\#; \Omega_n - \Omega_0)|\right\}\right] |Z_1^n| + o(1). \end{aligned}$$

Taking an expectation of the first term on the right over $Z_1, Z_2, \dots, Z_n \stackrel{\text{iid}}{\sim} P_0$, recalling that $w(\mathbb{H}_n^\# - \mathbb{H}_{n,0}^\#; \Omega_0) + w(\mathbb{H}_n^\#; \Omega_n - \Omega_0)$ is marginally $o_p(1)$, and applying the dominated convergence theorem shows that this nonnegative conditional expectation converges to zero in mean, and therefore also in probability. Hence, the above shows that $w(\mathbb{H}_n^\#; \Omega_n)$ converges weakly to $w(\mathbb{H}; \Omega_0)$, given $(Z_i)_{i=1}^\infty$, in probability. This implies that the $(1 - \alpha)$ quantile of $w(\mathbb{H}_n^\#; \Omega_n)$ conditional on Z_1^n , namely $\hat{\zeta}_n$, converges in probability to the $(1 - \alpha)$ -quantile of $w(\mathbb{H}; \Omega_0)$, namely $\zeta_{1-\alpha}$. \square

C.4 Proofs for Section 5

Proofs for Section 5.1

Recall that the (squared) Hilbert-Schmidt norm is defined as $\|\dot{\nu}_P^*\|_{\text{HS}}^2 := \sum_{k=1}^{\infty} \|\dot{\nu}_P^*(h_k)\|_{L^2(P)}^2$.

Lemma S10. *Suppose ν is pathwise differentiable at P with EIF ϕ_P . Then, $\|\phi_P\|_{L^2(P; \mathcal{H})}^2 = \|\dot{\nu}_P^*\|_{\text{HS}}^2$.*

Proof of Lemma S10. Suppose that ν has EIF ϕ_P . Using that (i) since ϕ_P is the EIF, it is P -a.s. true that $\dot{\nu}_P^*(h_k)(z) = \langle \phi_P(z), h_k \rangle_{\mathcal{H}}$ for all $k \in \mathbb{N}$, and (ii) for any $h \in \mathcal{H}$, $\|h\|_{\mathcal{H}}^2 = \sum_{k=1}^{\infty} \langle h, h_k \rangle_{\mathcal{H}}^2$, we see that

$$E_P \left[\sum_{k=1}^{\infty} \dot{\nu}_P^*(h_k)(Z)^2 \right] = E_P \left[\sum_{k=1}^{\infty} \langle \phi_P(Z), h_k \rangle_{\mathcal{H}}^2 \right] = E_P [\|\phi_P(Z)\|_{\mathcal{H}}^2] = \|\phi_P\|_{L^2(P; \mathcal{H})}^2. \quad (\text{S26})$$

□

In the following lemma, the little-oh and big-Omega notation both denote behavior as $n \rightarrow \infty$.

Lemma S11 (No tight, non-zero weak limit for a scaling of the the regularized one-step estimator when $\sum_{k=1}^{\infty} P_0 \dot{\nu}_0^*(h_k)^2 = +\infty$). *Suppose that ν is pathwise differentiable at P_0 . Let $(\beta_n)_{n=1}^{\infty}$ be an ℓ^2 -valued sequence that grows to $(1, 1, \dots)$ pointwise as $n \rightarrow \infty$ and let $(c_n)_{n=1}^{\infty}$ be a nonnegative real-valued sequence. All of the following hold:*

- (i) *if $c_n = o[n^{1/2}/\sigma_0(\beta_n)]$, then $c_n P_n \phi_0^{\beta_n} \xrightarrow{p} 0$;*
- (ii) *if $c_n = o(n^{1/2})$, then either $c_n P_n \phi_0^{\beta_n}$ does not converge weakly in \mathcal{H} to a tight random element or $c_n P_n \phi_0^{\beta_n} \xrightarrow{p} 0$;*
- (iii) *if $c_n = \Omega(n^{1/2})$ and $\sum_{k=1}^{\infty} P_0 \dot{\nu}_0^*(h_k)^2 = +\infty$, then $c_n P_n \phi_0^{\beta_n}$ does not converge weakly in \mathcal{H} to a tight random element.*

Before giving the proof, we note that the condition that $\sum_{k=1}^{\infty} P_0 \dot{\nu}_0^*(h_k)^2 = +\infty$ holds in all of the examples we exhibit in this work for which there does not exist an EIF. Moreover, if there does exist an EIF ϕ_0 , then Lemma S10 shows that $\sum_{k=1}^{\infty} P_0 \dot{\nu}_0^*(h_k)^2 < +\infty$ if and only if ϕ_0 is P_0 -Bochner square integrable.

Proof of Lemma S11. Let ν and $(\beta_n)_{n=1}^{\infty}$ be as in the statement of the lemma.

We first prove (i). Suppose that $c_n = o[n^{1/2}/\sigma_0(\beta_n)]$. By the definition of $\sigma_0^2(\beta_n)$, we have that

$$\begin{aligned} & E_{P_0^n} \|n^{1/2} P_n \phi_0^{\beta_n} / \sigma_0(\beta_n)\|_{\mathcal{H}}^2 \\ &= E_{P_0^n} \left[\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^{\infty} \beta_{n,k}^2 \dot{\nu}_0^*(h_k)(Z_i)^2 \right] / \sigma_0^2(\beta_n) = \sigma_0^2(\beta_n) / \sigma_0^2(\beta_n) = 1. \end{aligned}$$

As $c_n = o[n^{1/2}/\sigma_0(\beta_n)]$, this implies that $E_{P_0^n} \|c_n P_n \phi_0^{\beta_n}\|_{\mathcal{H}}^2 = o(1)$, which in turn implies that $\|c_n P_n \phi_0^{\beta_n}\|_{\mathcal{H}}^2$ is $o_p(1)$, as desired.

We now prove (ii). Suppose that $c_n = o(n^{1/2})$ and $c_n P_n \phi_0^{\beta_n}$ converges weakly in \mathcal{H} to a tight random element \mathbb{H}_0 . We will show that this can only be true if \mathbb{H}_0 is equal to the zero element of \mathcal{H} almost surely. By Theorem 1.8.4 of [van der Vaart and Wellner \(1996\)](#), $c_n P_n \phi_0^{\beta_n} \rightsquigarrow \mathbb{H}_0$ implies that $\langle c_n P_n \phi_0^{\beta_n}, h_k \rangle_{\mathcal{H}} \rightsquigarrow \langle \mathbb{H}_0, h_k \rangle_{\mathcal{H}}$ for all $k \in \mathbb{N}$. Moreover, since $\langle c_n P_n \phi_0^{\beta_n}, h_k \rangle_{\mathcal{H}} = c_n P_n \langle \phi_0^{\beta_n}, h_k \rangle_{\mathcal{H}} = \beta_{n,k} [c_n / n^{1/2}] [n^{1/2} P_n \dot{\nu}_0^*(h_k)]$, this shows that $\beta_{n,k} [c_n / n^{1/2}] [n^{1/2} P_n \dot{\nu}_0^*(h_k)] \rightsquigarrow \langle \mathbb{H}_0, h_k \rangle_{\mathcal{H}}$ for all $k \in \mathbb{N}$. Since $c_n / n^{1/2} = o(1)$, $\beta_{n,k} \xrightarrow{n \rightarrow \infty} 1$, and, by the central limit theorem, $n^{1/2} P_n \dot{\nu}_0^*(h_k) = O_p(1)$, it holds that $\beta_{n,k} [c_n / n^{1/2}] [n^{1/2} P_n \dot{\nu}_0^*(h_k)] \xrightarrow{p} 0$. As weak limits must share the same distribution, this shows that $\langle \mathbb{H}_0, h_k \rangle_{\mathcal{H}}$ is degenerate at zero for all k . Hence, \mathbb{H}_0 is almost surely equal to the zero element of \mathcal{H} .

We now prove (iii). It suffices to show that $c_n P_n \phi_0^{\beta_n}$ does not converge weakly to a tight random element when $c_n = n^{1/2}$. We argue this by contradiction. To this end, suppose that there exists a tight random element \mathbb{H}_0 such that $n^{1/2} P_n \phi_0^{\beta_n} \rightsquigarrow \mathbb{H}_0$. By Lemma 1.8.4 of [van der Vaart and Wellner \(1996\)](#), \mathbb{H}_0 is then such that $\langle n^{1/2} P_n \phi_0^{\beta_n}, h_k \rangle_{\mathcal{H}} \rightsquigarrow \langle \mathbb{H}_0, h_k \rangle_{\mathcal{H}}$ for all $k \in \mathbb{N}$. Combining this with the fact that $\langle n^{1/2} P_n \phi_0^{\beta_n}, h_k \rangle_{\mathcal{H}} = \beta_{n,k} n^{1/2} P_n \dot{\nu}_0^*(h_k)$, $\beta_{n,k} \xrightarrow{k \rightarrow \infty} 1$, and a univariate central limit theorem, this shows that \mathbb{H}_0 is such that $\langle \mathbb{H}_0, h_k \rangle_{\mathcal{H}} \sim N[0, \dot{\nu}_0^*(h_k)^2]$ for all $k \in \mathbb{N}$. Hence, \mathbb{H}_0 is a Gaussian random element. Also, by Fernique's theorem ([Fernique, 1970](#)), $\mathbb{H}_0 \in L^2(P_0; \mathcal{H})$. But $\|\mathbb{H}_0\|_{L^2(P_0; \mathcal{H})}^2 = E[\sum_{k=1}^{\infty} \langle \mathbb{H}_0, h_k \rangle_{\mathcal{H}}^2] = \sum_{k=1}^{\infty} P_0 \dot{\nu}_0^*(h_k)^2$, which is equal to $+\infty$ by assumption. Contradiction. \square

Proof of Theorem 5. By the definitions of $\bar{\nu}_n^{\beta_n}$, $\phi_0^{\beta_n}$, $\mathcal{R}_n^{j, \beta_n}$, and $\mathcal{D}_n^{j, \beta_n}$,

$$\begin{aligned} \bar{\nu}_n^{\beta_n} - \nu(P_0) - P_n \phi_0^{\beta_n} &= \frac{1}{2} \sum_{j=1}^2 \sum_{k=1}^{\infty} (1 - \beta_{n,k}) \langle \nu(\hat{P}_n^j) - \nu(P_0), h_k \rangle_{\mathcal{H}} h_k \\ &= \frac{1}{2} \sum_{j=1}^2 [\mathcal{R}_n^{j, \beta_n} + \mathcal{D}_n^{j, \beta_n}]. \end{aligned}$$

Taking an \mathcal{H} -norm of both sides, applying the triangle inequality on the right, and upper bounding averages by maxima yields

$$\left\| \bar{\nu}_n^{\beta_n} - \nu(P_0) - P_n \phi_0^{\beta_n} - \frac{1}{2} \sum_{j=1}^2 \mathcal{B}_n^{j, \beta_n} \right\|_{\mathcal{H}} \leq \max_j \|\mathcal{R}_n^{j, \beta_n}\|_{\mathcal{H}} + \max_j \|\mathcal{D}_n^{j, \beta_n}\|_{\mathcal{H}},$$

which bears resemblance to (21) but contains an extra bias term $\frac{1}{2} \sum_{j=1}^2 \mathcal{B}_n^{j, \beta_n}$. Plugging in the assumption that $\|\mathcal{R}_n^{j, \beta_n}\|_{\mathcal{H}}$ and $\|\mathcal{D}_n^{j, \beta_n}\|_{\mathcal{H}}$ are $O_p(\|\beta_n\|_{\ell^2}/n^{1/2})$ for each $j \in \{1, 2\}$ gives (25). Combining this with the assumption that $\mathcal{B}_n^{j, \beta_n} = O_p[\|\beta_n\|_{\ell^2}/n^{1/2}]$ for each $j \in \{1, 2\}$ and the fact that $P_n \phi_0^{\beta_n}$ is $O_p[\sigma_0(\beta_n)/n^{1/2}] = O_p[\|\beta_n\|_{\ell^2}/n^{1/2}]$ by Chebyshev's inequality ([Grenander, 1963](#)) and Lemma 1 then gives (27). \square

Proof of Lemma 4. This proof is similar to that of Lemma 3. Fix $\delta > 0$ and an ℓ^2 -valued sequence $(\beta_n)_{n=1}^{\infty}$ that is such that $\|\phi_n^{1, \beta_n} - \phi_0^{\beta_n}\|_{L^2(P_0; \mathcal{H})} = o_p(r_n)$ holds. We will show that

$$\lim_{n \rightarrow \infty} P_0^n \{ \|\mathcal{D}_n^{1, \beta_n}\|_{\mathcal{H}} > r_n n^{-1/2} \delta \} = 0. \quad (\text{S27})$$

As δ was arbitrary, this will show that $\mathcal{D}_n^{1, \beta_n} = o_p(r_n/n^{1/2})$. This will establish the stated result in the case where $j = 1$, and an analogous argument can be used to handle the case where $j = 2$.

Let $1_{\mathcal{E}_n}$ denote the indicator that the event \mathcal{E}_n that $\|\phi_n^{1, \beta_n} - \phi_0^{\beta_n}\|_{L^2(P_0; \mathcal{H})}^2 \leq r_n^2 \delta^2/2$ and let \mathcal{E}_n^c denote the complement of \mathcal{E}_n . We will leverage the following decomposition when showing (S27):

$$\begin{aligned} P_0^n \{ \|\mathcal{D}_n^{1, \beta_n}\|_{\mathcal{H}} > n^{-1/2} \delta \} &\leq P_0^n \left(\{ \|\mathcal{D}_n^{1, \beta_n}\|_{\mathcal{H}} > r_n n^{-1/2} \delta \} \cap \mathcal{E}_n \right) + P_0^n(\mathcal{E}_n^c) \\ &= E_0^n \left[1_{\mathcal{E}_n} P_0^n \left\{ \|\mathcal{D}_n^{1, \beta_n}\|_{\mathcal{H}} > r_n n^{-1/2} \delta \mid Z_1, \dots, Z_{n/2} \right\} \right] + o(1), \end{aligned} \quad (\text{S28})$$

where E_0^n denotes an expectation under sampling from the n -fold product measure P_0^n and $o(1)$ denotes a deterministic term that goes to zero as $n \rightarrow \infty$. The equality above holds by the law of total probability, the fact that \mathcal{E}_n is measurable with respect to the σ -field generated by $Z_1, \dots, Z_{n/2}$, and the assumption that $\|\phi_n^{1, \beta_n} - \phi_0^{\beta_n}\|_{L^2(P_0; \mathcal{H})} = o_p(r_n)$ implies that $P_0^n(\mathcal{E}_n^c) = o(1)$. To show (S27), the above shows that it suffices to show that the first term on the right-hand side is $o(1)$. To this end, note that Chebyshev's inequality for Hilbert-valued random variables ([Grenander, 1963](#)) and the bilinearity of inner products

shows that

$$\begin{aligned}
& 1_{\mathcal{E}_n} P_0^n \left\{ \|\mathcal{D}_n^{1,\beta_n}\|_{\mathcal{H}} > r_n n^{-1/2} \delta \mid Z_1, \dots, Z_{n/2} \right\} \\
& \leq 1_{\mathcal{E}_n} \frac{E_{P_0^n} [\|(P_n^1 - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})\|_{\mathcal{H}}^2 \mid Z_1, \dots, Z_{n/2}]}{r_n^2 n^{-1} \delta^2} \\
& = 1_{\mathcal{E}_n} \frac{(n/2)^{-1} E_{P_0^n} [P_n^1 \|(I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})\|_{\mathcal{H}}^2 \mid Z_1, \dots, Z_{n/2}]}{r_n^2 n^{-1} \delta^2} \\
& \quad + 1_{\mathcal{E}_n} \frac{4}{n^2} \sum_{i \neq j} \frac{E_{P_0^n} [\langle (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(Z_i), (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(Z_k) \rangle_{\mathcal{H}} \mid Z_1, \dots, Z_{n/2}]}{r_n^2 n^{-1} \delta^2},
\end{aligned} \tag{S29}$$

where $(I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})$ denotes the map $z \mapsto (\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(z) - P_0(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})$ and the sum is over $(i, j) \in \{n/2 + 1, \dots, n\}^2$ such that $i \neq j$. Each expectation in the summand on the right-hand side above is well-defined since $(\phi_n^{1,\beta_n} - \phi_0^{\beta_n}) \in L^2(P_0; \mathcal{H})$ implies that $(I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n}) \in L^2(P_0; \mathcal{H})$ as well. In fact, each expectation in the summand on the right-hand side is zero since, by the fact that $(I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n}) \in L^2(P_0; \mathcal{H})$ and Fubini's theorem,

$$\begin{aligned}
& E_{P_0^n} [\langle (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(Z_i), (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(Z_k) \rangle_{\mathcal{H}} \mid Z_1, \dots, Z_{n/2}] \\
& = \int \langle (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(z_1), (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(z_2) \rangle_{\mathcal{H}} P_0^2(dz_1, dz_2) \\
& = \iint \langle (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(z_1), (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(z_2) \rangle_{\mathcal{H}} P_0(dz_1) P_0(dz_2) \\
& = \int \left\langle \int (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(z_1) P_0(dz_1), (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(z_2) \right\rangle_{\mathcal{H}} P_0(dz_2) \\
& = \int \left\langle 0, (I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})(z_2) \right\rangle_{\mathcal{H}} P_0(dz_2) = 0.
\end{aligned}$$

Returning to (S29) and simplifying the first term on the right-hand side of that expression, this shows that

$$1_{\mathcal{E}_n} P_0^n \left\{ \|\mathcal{D}_n^{1,\beta_n}\|_{\mathcal{H}} > r_n n^{-1/2} \delta \mid Z_1, \dots, Z_{n/2} \right\} \leq 1_{\mathcal{E}_n} \frac{2 \|(I - P_0)(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})\|_{L^2(P_0; \mathcal{H})}^2}{r_n^2 \delta^2}.$$

Using that $P_0(\phi_n^{1,\beta_n} - \phi_0^{\beta_n})$ is a minimizer over $h \in \mathcal{H}$ of $\|\phi_n^{1,\beta_n} - \phi_0^{\beta_n} - h\|_{L^2(P_0; \mathcal{H})}^2$ and subsequently leveraging the definition of the event \mathcal{E}_n , this shows that

$$\begin{aligned}
1_{\mathcal{E}_n} P_0^n \left\{ \|\mathcal{D}_n^{1,\beta_n}\|_{\mathcal{H}} > r_n n^{-1/2} \delta \mid Z_1, \dots, Z_{n/2} \right\} & \leq 1_{\mathcal{E}_n} \frac{2 \|\phi_n^{1,\beta_n} - \phi_0^{\beta_n}\|_{L^2(P_0; \mathcal{H})}^2}{r_n^2 \delta^2} \\
& \leq \min \left\{ 1, \frac{2 \|\phi_n^{1,\beta_n} - \phi_0^{\beta_n}\|_{L^2(P_0; \mathcal{H})}^2}{r_n^2 \delta^2} \right\}.
\end{aligned}$$

Note that the right-hand side above is no larger than 1. Taking an expectation of both sides over $Z_1, \dots, Z_{n/2} \stackrel{\text{iid}}{\sim} P_0$ and recalling that $\|\phi_n^{1,\beta_n} - \phi_0^{\beta_n}\|_{L^2(P_0; \mathcal{H})} = o_p(r_n)$, the dominated convergence theorem shows that the first term on the right-hand side of (S28) is $o(1)$. This completes the proof. \square

Proof of Lemma 5. Since $\nu(P), \nu(P_0) \in \mathcal{H}$ and $\phi_P^\beta \in L^2(P_0; \mathcal{H})$, it holds that $\mathcal{R}_P^\beta \in \mathcal{H}$. We begin by showing that, for any $\ell \in \mathbb{N}$,

$$\left\langle \mathcal{R}_P^\beta, h_\ell \right\rangle_{\mathcal{H}} = \beta_\ell [\langle \nu(P) - \nu(P_0), h_\ell \rangle_{\mathcal{H}} + P_0 \nu_P^*(h_\ell)]. \tag{S30}$$

As ℓ was arbitrary and $(h_k)_{k=1}^\infty$ is an orthonormal basis of \mathcal{H} , this will then show that

$$\mathcal{R}_P^\beta = \sum_{k=1}^\infty \beta_k [\langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} + P_0 \dot{\nu}_P^*(h_k)] h_k. \quad (\text{S31})$$

We now establish (S30) for a fixed $\ell \in \mathbb{N}$. Note that

$$\begin{aligned} & \left\langle \mathcal{R}_P^\beta, h_\ell \right\rangle_{\mathcal{H}} \\ &:= \left\langle \nu(P) - \nu(P_0) + P_0 \phi_P^\beta - \sum_{k=1}^\infty (1 - \beta_k) \langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} h_k, h_\ell \right\rangle_{\mathcal{H}} \\ &= \left\langle \sum_{k=1}^\infty \langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} h_k + P_0 \phi_P^\beta - \sum_{k=1}^\infty (1 - \beta_k) \langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} h_k, h_\ell \right\rangle_{\mathcal{H}} \\ &= \left\langle \sum_{k=1}^\infty \beta_k \langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}} h_k + P_0 \phi_P^\beta, h_\ell \right\rangle_{\mathcal{H}} \\ &= \beta_\ell \langle \nu(P) - \nu(P_0), h_\ell \rangle_{\mathcal{H}} + \left\langle P_0 \phi_P^\beta, h_\ell \right\rangle_{\mathcal{H}}. \end{aligned}$$

It remains to show that $\langle P_0 \phi_P^\beta, h_\ell \rangle_{\mathcal{H}} = \beta_\ell P_0 \dot{\nu}_P^*(h_\ell)$. To see that this holds, note that, since ν is pathwise differentiable at P , Lemma 1 ensures that $\phi_P^\beta(z)$ is the Riesz representation of $r_P^\beta(\cdot)(z)$ on a set \mathcal{Z}^β of P -probability one. Since $P_0 \ll P$, \mathcal{Z}^β has P_0 -probability one as well. Hence, $\int \langle \phi_P^\beta(z), h_\ell \rangle_{\mathcal{H}} P_0(dz) = \int r_P^\beta(h_\ell)(z) P_0(dz) = P_0 r_P^\beta(h_\ell)$. Furthermore, since $\phi_P^\beta \in L^2(P_0; \mathcal{H})$, $\int \langle \phi_P^\beta(z), h_\ell \rangle_{\mathcal{H}} P_0(dz) = \langle P_0 \phi_P^\beta, h_\ell \rangle_{\mathcal{H}}$, and so $\langle P_0 \phi_P^\beta, h_\ell \rangle_{\mathcal{H}} = P_0 r_P^\beta(h_\ell)$. Plugging in the definition of r_P^β shows that $\langle P_0 \phi_P^\beta, h_\ell \rangle_{\mathcal{H}} = \beta_\ell P_0 \dot{\nu}_P^*(h_\ell)$, as desired. This establishes (S30), which in turn establishes (S31). Using the form of \mathcal{R}_P^β given in (S31) establishes the equality in the statement of the lemma. The inequality follows by Cauchy-Schwarz. \square

Proof of Lemma 6. We have that

$$\begin{aligned} & \|\mathcal{B}_P^\beta\|_{\mathcal{H}}^2 \\ &= \sum_{k=1}^\infty (1 - \beta_k)^2 \langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}}^2 = \sum_{k=1}^\infty k^{-2u} (1 - \beta_k)^2 k^{2u} \langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}}^2 \\ &\leq \left[\sup_{k \in \mathbb{N}} \frac{(1 - \beta_k)^2}{k^{2u}} \right] \sum_{k=1}^\infty k^{2u} \langle \nu(P) - \nu(P_0), h_k \rangle_{\mathcal{H}}^2 = \left[\sup_{k \in \mathbb{N}} \frac{(1 - \beta_k)^2}{k^{2u}} \right] \|\nu(P) - \nu(P_0)\|_u^2. \end{aligned}$$

Taking a square root of both sides above gives the inequality from the lemma statement for general β . In the special case where $\beta_k = 1$ for all $k \leq K$ and $\beta_k = 0$ for all $k > K$, (28) follows by plugging this value of β and then applying the triangle inequality. \square

C.4.1 Proofs for Section 5.2

Proof of Lemma 7. In this proof, we will use that $\Gamma_\beta : \mathcal{H} \rightarrow \mathcal{H}$ is a linear mapping, and also that $\|\Gamma_\beta(h)\|_{\mathcal{H}} \leq \|h\|_{\mathcal{H}}$ for all h .

We now show that $\Gamma_\beta \circ \nu$ is pathwise differentiable with local parameter $\Gamma_\beta \circ \dot{\nu}_P$. To see that this holds, fix a quadratic mean differentiable submodel $\{P_\epsilon : \epsilon \in [0, \delta)\} \in \mathcal{P}(P, \mathcal{P}, s)$. Note that

$$\begin{aligned} \|\Gamma_\beta \circ \nu(P_\epsilon) - \Gamma_\beta \circ \nu(P) - \epsilon \Gamma_\beta \circ \dot{\nu}_P(s)\|_{\mathcal{H}} &= \|\Gamma_\beta \circ [\nu(P_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s)]\|_{\mathcal{H}} \\ &\leq \|\nu(P_\epsilon) - \nu(P) - \epsilon \dot{\nu}_P(s)\|_{\mathcal{H}} = o(\epsilon), \end{aligned}$$

where the final equality holds by the pathwise differentiability of ν . Hence, $\Gamma_\beta \circ \nu$ is pathwise differentiable

with local parameter $\Gamma_\beta \circ \dot{\nu}_P$. The efficient influence operator is equal to r_P^β , where this quantity is as defined above Lemma 1. Indeed, for any $s \in \dot{\mathcal{P}}_P$ and $h \in \mathcal{H}$,

$$\begin{aligned} \left\langle s, r_P^\beta(h) \right\rangle_{L^2(P)} &= \left\langle s, \sum_{k=1}^{\infty} \beta_k \langle h, h_k \rangle_{\mathcal{H}} \dot{\nu}_P^*(h_k) \right\rangle_{L^2(P)} = \sum_{k=1}^{\infty} \beta_k \langle h, h_k \rangle_{\mathcal{H}} \langle s, \dot{\nu}_P^*(h_k) \rangle_{L^2(P)} \\ &= \sum_{k=1}^{\infty} \beta_k \langle h, h_k \rangle_{\mathcal{H}} \langle \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} = \left\langle h, \sum_{k=1}^{\infty} \beta_k \langle \dot{\nu}_P(s), h_k \rangle_{\mathcal{H}} h_k \right\rangle_{\mathcal{H}} \\ &= \langle h, \Gamma_\beta \circ \dot{\nu}_P(s) \rangle_{\mathcal{H}}, \end{aligned}$$

where above we have used the linearity and continuity of inner products and the definition of the efficient influence operator $\dot{\nu}_P^*$ of ν . By Lemma 1, $r_P^\beta(\cdot)(z)$ is P -almost surely a bounded linear operator with Riesz representation $\phi_P^\beta \in L^2(P; \mathcal{H})$, and therefore ϕ_P^β is the EIF of $\Gamma_\beta \circ \nu$. \square

The following is a consequence of Theorem 2, specialized to the case where the pathwise differentiable parameter of interest takes the form $\nu^\beta := \Gamma_\beta \circ \nu$. Below $\mathcal{R}_n^{j,\beta}$ and $\mathcal{D}_n^{j,\beta}$ are the regularized remainder and drift terms for the β -regularized one-step estimator $\bar{\nu}_n^\beta$ of $\nu(P_0)$, as defined in Section 5.1.

Corollary S1 (Asymptotic linearity of $\tilde{\nu}_n^\beta$). *Fix $\beta \in \ell_*^2$. Suppose that ν is pathwise differentiable at P_0 and that $\mathcal{R}_n^{j,\beta}$ and $\mathcal{D}_n^{j,\beta}$ are both $o_p(n^{-1/2})$ for $j \in \{1, 2\}$. Under these conditions,*

$$\tilde{\nu}_n^\beta - \nu^\beta(P_0) = \frac{1}{n} \sum_{i=1}^n \phi_0^\beta(Z_i) + o_p(n^{-1/2}), \quad (\text{S32})$$

$\tilde{\nu}_n^\beta$ is regular, and $n^{1/2}[\tilde{\nu}_n^\beta - \nu^\beta(P_0)] \rightsquigarrow \mathbb{H}$, where \mathbb{H} is a tight \mathcal{H} -valued Gaussian random variable that is such that, for each $h \in \mathcal{H}$, the marginal distribution $\langle \mathbb{H}, h \rangle_{\mathcal{H}}$ follows a $N(0, E_0[\langle \phi_0^\beta(Z), h \rangle_{\mathcal{H}}^2])$ distribution.

Since the above imposes conditions on the regularized remainder and drift terms for the regularized one-step estimator $\bar{\nu}_n^\beta$, any analysis that is performed to bound these terms when studying $\bar{\nu}_n^\beta$ can also be used to bound these terms when studying $\tilde{\nu}_n^\beta$. In particular, Lemmas 4 and 5 can be used to study these terms.

Proof of Corollary S1. We establish that the conditions of Theorem 2 are satisfied. By Lemma 7, the EIF of ν^β is equal to $\phi_0^\beta \in L^2(P_0; \mathcal{H})$ at P_0 and $\phi_n^{j,\beta}$ at \hat{P}_n^j , $j \in \{1, 2\}$. Hence, for each $j \in \{1, 2\}$, the regularized drift term $\mathcal{D}_n^{j,\beta} := (\hat{P}_n^j - P_0)(\phi_n^{j,\beta} - \phi_0^\beta)$ for the β -regularized one-step estimator $\bar{\nu}_n^\beta$ of $\nu(P_0)$ is identical to the drift term for the one-step estimator $\tilde{\nu}_n^\beta$ of $\nu^\beta(P_0)$. Moreover, since

$$\begin{aligned} \mathcal{R}_n^{j,\beta} &:= \nu(\hat{P}_n^j) - \nu(P_0) + P_0 \phi_n^{j,\beta} - \sum_{k=1}^{\infty} (1 - \beta_k) \langle \nu(\hat{P}_n^j) - \nu(P_0), h_k \rangle_{\mathcal{H}} h_k \\ &= \nu^\beta(\hat{P}_n^j) - \nu^\beta(P_0) + P_0 \phi_n^{j,\beta}, \end{aligned}$$

the regularized remainder term $\mathcal{R}_n^{j,\beta}$ for the β -regularized one-step estimator $\bar{\nu}_n^\beta$ of $\nu(P_0)$ is also identical to the remainder term for the one-step estimator $\tilde{\nu}_n^\beta$ of $\nu^\beta(P_0)$. As we have assumed that both $\mathcal{D}_n^{j,\beta}$ and $\mathcal{R}_n^{j,\beta}$ are $o_p(n^{-1/2})$, $j \in \{1, 2\}$, Theorem 2 implies all the claims in the statement of this corollary. \square

Proof of Theorem 6. By Lemma 7, ϕ_0^β is the EIF of ν^β at P_0 . By the same arguments used to establish Theorem 2, $\tilde{\nu}_n^\beta$ is a regular estimator of $\nu^\beta(P_0)$, and, in particular, $n^{1/2}[\tilde{\nu}_n^\beta - \nu^\beta(P_\epsilon)] \rightsquigarrow \mathbb{H}^\beta$ under the sampling of n iid draws from $P_{\epsilon=n^{-1/2}}$. Combining this with the pathwise differentiability of ν^β , this shows that $n^{1/2}[\tilde{\nu}_n^\beta - \nu^\beta(P_0)] \rightsquigarrow \mathbb{H}^\beta + \dot{\nu}_0^\beta(s)$, where $\dot{\nu}_0^\beta := \Gamma_\beta \circ \dot{\nu}_0$ is the local parameter of ν^β at P_0 .

By similar arguments to those used in the proof of Theorem 3, $\|\mathbb{H}^\beta + \dot{\nu}_0^\beta(s)\|_{\mathcal{H}}^2$ is a continuous random variable. Combining this with the fact that $\widehat{\zeta}_n \rightarrow \zeta_{1-\alpha}$ shows that

$$P_{\epsilon=n^{-1/2}}^n \left\{ h_0 \notin \Gamma_\beta^{-1}[\mathcal{C}_n^\beta(\widehat{\zeta}_n)] \right\} \xrightarrow{n \rightarrow \infty} \Pr \left\{ \|\mathbb{H}^\beta + \dot{\nu}_0^\beta(s)\|_{\mathcal{H}}^2 > \zeta_{1-\alpha} \right\}.$$

By Corollary 2 of [Lewandowski et al. \(1995\)](#), the definition of $\zeta_{1-\alpha}$, and the fact that $\|\dot{\nu}_0(s)\|_{\mathcal{H}} > 0$ and $\beta > 0$ entrywise together imply that $\|\dot{\nu}_0^\beta(s)\|_{\mathcal{H}} > 0$, the right-hand side above is strictly larger than α .

To see that $\nu(P_{\epsilon=n^{-1/2}})$ is an $n^{-1/2}$ -rate local alternative, note that

$$\begin{aligned} \|\nu(P_{\epsilon=n^{-1/2}}) - h_0\|_{\mathcal{H}} &= \|\nu(P_{\epsilon=n^{-1/2}}) - \nu(P_0)\|_{\mathcal{H}} \\ &= n^{-1/2} \|\dot{\nu}_0(s)\|_{\mathcal{H}} + o(n^{-1/2}) = O(n^{-1/2}). \end{aligned}$$

Above we used that $\nu(P_0) = h_0$, $\{P_\epsilon : \epsilon\}$ is quadratic mean differentiable, and $\dot{\nu}_0$ is the local parameter of ν at P_0 . \square

D A regularized inverse covariance operator and a consistent estimator thereof

Let $\Sigma_0 : h \mapsto E[\langle h, \mathbb{H} \rangle_{\mathcal{H}} \mathbb{H}]$ denote the covariance operator of the Gaussian random element \mathbb{H} from Theorem 2. In this appendix, we study the regularized inverse $\Omega_0 = [(1-\lambda)\Sigma_0 + \lambda I]^{-1}$, where $\lambda > 0$ and I denotes the identity operator on \mathcal{H} . Though it would be interesting to study the behavior of our confidence set in cases where λ shrinks to zero with sample size, doing so may be challenging since the inverse covariance operator Σ_0^{-1} may not exist and, even if it does, it will generally be unbounded, which will complicate the use of the continuous mapping theorem that we use to justify the proof of the asymptotic validity of our confidence set (Theorem 3). Hence, while studying the case where λ shrinks to zero slowly with sample size is an interesting area for future work, here we focus on the case where λ is a fixed constant that does not depend on sample size.

The regularized inverse of interest writes as $\Omega_0 = f_\lambda(\Sigma_0)$, where, for a positive semidefinite linear operator $\Sigma : \mathcal{H} \rightarrow \mathbb{H}$, $f_\lambda(\Sigma) := [(1-\lambda)\Sigma + \lambda I]^{-1}$. The operator f_λ can be seen to be Lipschitz continuous relative to the operator norm with Lipschitz constant $(1-\lambda)/\lambda^2$, which holds since, for positive semidefinite Σ_1 and Σ_2 ,

$$\begin{aligned} &\|f_\lambda(\Sigma_1) - f_\lambda(\Sigma_2)\|_{\text{op}} \\ &= \left\| [(1-\lambda)\Sigma_1 + \lambda I]^{-1} \circ \{[(1-\lambda)\Sigma_2 + \lambda I] - [(1-\lambda)\Sigma_1 + \lambda I]\} \circ [(1-\lambda)\Sigma_2 + \lambda I]^{-1} \right\|_{\text{op}} \\ &= (1-\lambda) \left\| [(1-\lambda)\Sigma_1 + \lambda I]^{-1} \circ (\Sigma_1 - \Sigma_2) \circ [(1-\lambda)\Sigma_2 + \lambda I]^{-1} \right\|_{\text{op}} \\ &\leq (1-\lambda) \|(1-\lambda)\Sigma_1 + \lambda I\|_{\text{op}} \|\Sigma_1 - \Sigma_2\|_{\text{op}} \|(1-\lambda)\Sigma_2 + \lambda I\|_{\text{op}} \\ &\leq (1-\lambda) \lambda^{-2} \|\Sigma_1 - \Sigma_2\|_{\text{op}}. \end{aligned}$$

By the continuous mapping theorem, an operator-norm-consistent estimator Ω_n of Ω_0 — that is, one for which $\|\Omega_n - \Omega_0\|_{\text{op}} = o_p(1)$ — is thus given by $\Omega_n = f_\lambda(\Sigma_n)$, where Σ_n is any operator-norm consistent estimator of Σ_0 . The following lemma shows that one such estimator is given by $h \mapsto \frac{1}{2} \sum_{j=1}^2 E_{P_n^j} [\langle h, \phi_n^j(Z) \rangle_{\mathcal{H}} \phi_n^j(Z)]$.

Lemma S12. *Fix $\lambda > 0$. Suppose that $\|\phi_0\|_{L^2(P_0; \mathcal{H})} < \infty$ and $\|\phi_n^j - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$ for each $j \in \{1, 2\}$. If $\Sigma_n : h \mapsto \frac{1}{2} \sum_{j=1}^2 E_{P_n^j} [\langle h, \phi_n^j(Z) \rangle_{\mathcal{H}} \phi_n^j(Z)]$, then $\|\Sigma_n - \Sigma_0\|_{\text{op}} = o_p(1)$.*

In what follows, for a function $\phi : \mathcal{Z} \rightarrow \mathcal{H}$, we let $\langle h, \phi \rangle_{\mathcal{H}}$ denote the map $z \mapsto \langle h, \phi(z) \rangle_{\mathcal{H}}$. We also recall

that \mathcal{H}_1 denotes the unit ball of \mathcal{H} . We give the proof of the above result after we prove the following supporting lemma.

Lemma S13. *In the setting of Lemma S12, $\{z \mapsto \langle h, \phi_0(z) \rangle_{\mathcal{H}}^2 : h \in \mathcal{H}_1\}$ is P_0 -Glivenko-Cantelli. Hence, $\sup_{h \in \mathcal{H}_1} (P_n^1 - P_0) \langle h, \phi_0 \rangle_{\mathcal{H}}^2 = o_p(1)$.*

Proof of Lemma S13. In what follows we let $\mathcal{F} := \{\langle h, \phi_0 \rangle_{\mathcal{H}}^2 : h \in \mathcal{H}_1\}$. We will show that this collection of functions is P_0 -Glivenko-Cantelli, which is the first result in the statement of the lemma. Combining this with the fact that P_n^1 is the empirical distribution of an iid sample from P_0 will then give the second result. To establish that \mathcal{F} is Glivenko-Cantelli, we will show that the conditions of Theorem 2.4.3 in [van der Vaart and Wellner \(1996\)](#) are satisfied. These conditions follow from \mathcal{F} having a P_0 -integrable envelope function F and, moreover, satisfying an appropriate covering number condition — we will define this condition in the next paragraph. Before doing so, we note that \mathcal{F} has P_0 -integrable envelope function $F(z) := \|\phi_0(z)\|_{\mathcal{H}}^2$. To see that this function is indeed an envelope of \mathcal{F} , note that, for any $h \in \mathcal{H}_1$, the Cauchy-Schwarz inequality shows that $\langle h, \phi_0(z) \rangle_{\mathcal{H}}^2 \leq \|\phi_0(z)\|_{\mathcal{H}}^2 = F(z)$. To see that F is P_0 -integrable, note that $P_0 F = \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$, which is finite by an assumption of Lemma S12.

In the remainder of this proof, we will establish a covering number condition on \mathcal{F} that implies the covering number condition from Theorem 2.4.3 in [van der Vaart and Wellner \(1996\)](#). In particular, in what follows we will show that, for any $\epsilon > 0$, there exists an $N \in \mathbb{N}$ such that, with probability tending to one, the $L^1(P_n^1)$ -covering number of \mathcal{F} is no more than N ; here we recall that, for fixed $\epsilon > 0$, the corresponding $L^1(P_n^1)$ covering number of \mathcal{F} denotes the size of the minimal ϵ -cover of \mathcal{F} relative to the $L^1(P_n^1)$ metric. Problem 2.4.2 in [van der Vaart and Wellner \(1996\)](#) justifies why this condition suffices to establish the covering number condition in Theorem 2.4.3 of that reference.

Fix $\epsilon > 0$ and an orthonormal basis $(h_k)_{k=1}^{\infty}$ of \mathcal{H} . By the monotone convergence theorem and the P_0 -Bochner square integrability of ϕ_0 ,

$$\lim_{K' \rightarrow \infty} P_0 \left[\sum_{k=1}^{K'} \langle h_k, \phi_0 \rangle_{\mathcal{H}}^2 \right] = P_0 \left[\sum_{k=1}^{\infty} \langle h_k, \phi_0 \rangle_{\mathcal{H}}^2 \right] = \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 < \infty.$$

Hence, there exists a $K < \infty$ such that $P_0 \left[\sum_{k=1}^K \langle h_k, \phi_0 \rangle_{\mathcal{H}}^2 \right] > \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 - \epsilon/8$, and so, for this K , $P_0 \left[\sum_{k=K+1}^{\infty} \langle h_k, \phi_0 \rangle_{\mathcal{H}}^2 \right] \leq \epsilon/8$. By the weak law of large numbers, we further have that

$$P_n^1 \left[\sum_{k=K+1}^{\infty} \langle h_k, \phi_0 \rangle_{\mathcal{H}}^2 \right] = P_0 \left[\sum_{k=K+1}^{\infty} \langle h_k, \phi_0 \rangle_{\mathcal{H}}^2 \right] + o_p(1).$$

Hereafter we work on the event \mathcal{E}_n where (i) the $o_p(1)$ term above is less than $\epsilon/8$, so that the left-hand side above is no more than $\epsilon/4$, and (ii) $\|\phi_0\|_{L^2(P_n^1; \mathcal{H})}^2 \leq \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 + \epsilon$; note that \mathcal{E}_n holds with probability tending to one as $n \rightarrow \infty$. We now show that there exists a fixed subset $\tilde{\mathcal{H}}_1$ of \mathcal{H}_1 such that, on this event, $\mathcal{F}_\epsilon := \{\langle h, \phi_0 \rangle_{\mathcal{H}}^2 : h \in \tilde{\mathcal{H}}_1\}$ is an ϵ -cover of \mathcal{F} . In particular, we take $\tilde{\mathcal{H}}_1$ to be a finite δ -cover of the finite-dimensional subset $\tilde{\mathcal{H}}_1 := \mathcal{H}_1 \cap \text{span}\{h_1, \dots, h_K\}$ of \mathcal{H}_1 relative to the \mathcal{H} -norm, where $\delta := \epsilon/[4(\|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 + \epsilon)]$. Such a finite δ -cover is guaranteed to exist because the unit ball in a finite-dimensional Hilbert space is necessarily totally bounded in the norm topology. To see that \mathcal{F}_ϵ is indeed an ϵ -cover of \mathcal{F} , fix $h \in \mathcal{H}$ and let $\tilde{h} \in \tilde{\mathcal{H}}_1$ be such that $\|\pi_K h - \tilde{h}\|_{\mathcal{H}} \leq \delta$, where $\pi_K h := \Pi_{\mathcal{H}}(h \mid \text{span}\{h_1, \dots, h_K\})$. Observe that

$$\begin{aligned} & \left\| \langle h, \phi_0 \rangle_{\mathcal{H}}^2 - \langle \tilde{h}, \phi_0 \rangle_{\mathcal{H}}^2 \right\|_{L^1(P_n^1)} \\ & \leq \left\| \langle h, \phi_0 \rangle_{\mathcal{H}}^2 - \langle \pi_K h, \phi_0 \rangle_{\mathcal{H}}^2 \right\|_{L^1(P_n^1)} + \left\| \langle \pi_K h, \phi_0 \rangle_{\mathcal{H}}^2 - \langle \tilde{h}, \phi_0 \rangle_{\mathcal{H}}^2 \right\|_{L^1(P_n^1)}. \end{aligned} \tag{S33}$$

We now show each of the two terms on the right-hand side is no more than $\epsilon/2$. For the first term, we let $\pi_K \phi_0(z) := \Pi_{\mathcal{H}}[\phi_0(z) \mid \text{span}\{h_1, \dots, h_K\}]$ and note that

$$\begin{aligned}
& \|\langle h, \phi_0 \rangle_{\mathcal{H}}^2 - \langle \pi_K h, \phi_0 \rangle_{\mathcal{H}}^2\|_{L^1(P_n^1)} \\
&= \int |\langle h + \pi_K h, \phi_0(z) \rangle_{\mathcal{H}} \langle h - \pi_K h, \phi_0(z) \rangle_{\mathcal{H}}| P_n^1(dz) \\
&\leq \left[\int \langle h + \pi_K h, \phi_0(z) \rangle_{\mathcal{H}}^2 P_n^1(dz) \right]^{1/2} \left[\int \langle h - \pi_K h, \phi_0(z) \rangle_{\mathcal{H}}^2 P_n^1(dz) \right]^{1/2} \\
&= \left[\int \langle h + \pi_K h, \phi_0(z) \rangle_{\mathcal{H}}^2 P_n^1(dz) \right]^{1/2} \left[\int \langle h - \pi_K h, \phi_0(z) - \pi_K \phi_0(z) \rangle_{\mathcal{H}}^2 P_n^1(dz) \right]^{1/2} \\
&\leq \|h + \pi_K h\|_{\mathcal{H}} \|h - \pi_K h\|_{\mathcal{H}} \|\phi_0 - \pi_K \phi_0\|_{L^2(P_n^1; \mathcal{H})}^2.
\end{aligned}$$

The first equality holds by the definition of the $L^1(P_n^1)$ norm, both inequalities hold by Cauchy-Schwarz, and the second equality holds because $h - \pi_K h$ is orthogonal to $\text{span}\{h_1, \dots, h_K\}$. Now, since $h \in \mathcal{H}_1$, the triangle inequality and the fact that orthogonal projections cannot increase length show that $\|h + \pi_K h\|_{\mathcal{H}} \|h - \pi_K h\|_{\mathcal{H}} \leq 2$. Furthermore, by the choice of K and the fact that we are working on the event \mathcal{E}_n , $\|\phi_0 - \pi_K \phi_0\|_{L^2(P_n^1; \mathcal{H})}^2 \leq \epsilon/4$. Hence, the first term on the right-hand side of (S33) is no more than $\epsilon/2$. For the second term in (S33), two consecutive applications of the Cauchy-Schwarz inequality yield that

$$\begin{aligned}
& \|\langle \pi_K h, \phi_0 \rangle_{\mathcal{H}}^2 - \langle \tilde{h}, \phi_0 \rangle_{\mathcal{H}}^2\|_{L^1(P_n^1)} \\
&= \int |\langle \pi_K h + \tilde{h}, \phi_0(z) \rangle_{\mathcal{H}} \langle \pi_K h - \tilde{h}, \phi_0(z) \rangle_{\mathcal{H}}| P_n^1(dz) \\
&\leq \left[\int \langle \pi_K h + \tilde{h}, \phi_0(z) \rangle_{\mathcal{H}}^2 P_n^1(dz) \right]^{1/2} \left[\int \langle \pi_K h - \tilde{h}, \phi_0(z) \rangle_{\mathcal{H}}^2 P_n^1(dz) \right]^{1/2} \\
&\leq \|\pi_K h + \tilde{h}\|_{\mathcal{H}} \|\pi_K h - \tilde{h}\|_{\mathcal{H}} \|\phi_0\|_{L^2(P_n^1; \mathcal{H})}^2.
\end{aligned}$$

Now, by the triangle inequality and the fact that $\pi_K h$ and \tilde{h} belong to \mathcal{H}_1 , $\|\pi_K h + \tilde{h}\|_{\mathcal{H}} \leq 2$. Moreover, because we are working on the event \mathcal{E}_n , $\|\phi_0\|_{L^2(P_n^1; \mathcal{H})}^2 \leq \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 + \epsilon$. Combining these bounds with the fact that $\|\pi_K h - \tilde{h}\|_{\mathcal{H}} \leq \delta := \epsilon/[4(\|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 + \epsilon)]$ gives that $\|\langle \pi_K h, \phi_0 \rangle_{\mathcal{H}}^2 - \langle \tilde{h}, \phi_0 \rangle_{\mathcal{H}}^2\|_{L^1(P_n^1)} \leq \epsilon/2$. Returning to (S33), this shows that $\|\langle h, \phi_0 \rangle_{\mathcal{H}}^2 - \langle \tilde{h}, \phi_0 \rangle_{\mathcal{H}}^2\|_{L^1(P_n^1)} \leq \epsilon$. As $h \in \mathcal{H}_1$ was arbitrary and $\langle \tilde{h}, \phi_0 \rangle_{\mathcal{H}}^2 \in \mathcal{F}_\epsilon$, this shows that \mathcal{F}_ϵ is an ϵ -cover of \mathcal{H}_1 on the event \mathcal{E}_n . Since \mathcal{F}_ϵ contains finitely many functions, we can invoke Theorem 2.4.3 of [van der Vaart and Wellner \(1996\)](#) to show that \mathcal{F} is a P_0 -Glivenko Cantelli class. \square

We conclude this appendix with the proof of Lemma S12.

Proof of Lemma S12. For $j \in \{1, 2\}$, let $\Sigma_n^j(h) := E_{P_n^j}[\langle h, \phi_n^j(Z) \rangle_{\mathcal{H}} \phi_n^j(Z)]$. By the triangle inequality,

$$\|\Sigma_n - \Sigma_0\|_{\text{op}} = \left\| \frac{1}{2} \sum_{j=1}^2 \Sigma_n^j - \Sigma_0 \right\|_{\text{op}} \leq \frac{1}{2} \sum_{j=1}^2 \|\Sigma_n^j - \Sigma_0\|_{\text{op}}.$$

Hence, it suffices to show that $\|\Sigma_n^j - \Sigma_0\|_{\text{op}} = o_p(1)$. We show this for the case where $j = 1$, and the case where $j = 2$ follows by analogous arguments.

Because $\Sigma_n^1 - \Sigma_0$ is a positive, self-adjoint operator, it holds that

$$\|\Sigma_n^1 - \Sigma_0\|_{\text{op}} = \sup_{h \in \mathcal{H}_1} \langle \Sigma_n^1(h) - \Sigma_0(h), h \rangle_{\mathcal{H}}. \tag{S34}$$

We will bound the right-hand side above in what follows. To do this, we will use that, for any $h \in \mathcal{H}_1$,

$$\begin{aligned}\langle \Sigma_n^1(h) - \Sigma_0(h), h \rangle_{\mathcal{H}} &= \langle P_n^1 \langle h, \phi_n^1 \rangle_{\mathcal{H}} \phi_n^1 - P_0 \langle h, \phi_0 \rangle_{\mathcal{H}} \phi_0, h \rangle_{\mathcal{H}} \\ &= \langle P_n^1 [\langle h, \phi_n^1 \rangle_{\mathcal{H}} \phi_n^1 - \langle h, \phi_0 \rangle_{\mathcal{H}} \phi_0] + (P_n^1 - P_0) \langle h, \phi_0 \rangle_{\mathcal{H}} \phi_0, h \rangle_{\mathcal{H}} \\ &= P_n^1 [\langle h, \phi_n^1 \rangle_{\mathcal{H}}^2 - \langle h, \phi_0 \rangle_{\mathcal{H}}^2] + (P_n^1 - P_0) \langle h, \phi_0 \rangle_{\mathcal{H}}^2 \\ &= P_n^1 \langle h, \phi_n^1 - \phi_0 \rangle_{\mathcal{H}} \langle h, \phi_n^1 + \phi_0 \rangle_{\mathcal{H}} + (P_n^1 - P_0) \langle h, \phi_0 \rangle_{\mathcal{H}}^2.\end{aligned}$$

Applying the triangle and Cauchy-Schwarz inequalities to the above and combining the result with (S34) shows that

$$\begin{aligned}\|\Sigma_n^1 - \Sigma_0\|_{\text{op}} &\leq \left[\sup_{h \in \mathcal{H}_1} P_n^1 \langle h, \phi_n^1 - \phi_0 \rangle_{\mathcal{H}}^2 \right]^{1/2} \left[\sup_{h \in \mathcal{H}_1} P_n^1 \langle h, \phi_n^1 + \phi_0 \rangle_{\mathcal{H}}^2 \right]^{1/2} \\ &\quad + \sup_{h \in \mathcal{H}_1} (P_n^1 - P_0) \langle h, \phi_0 \rangle_{\mathcal{H}}^2.\end{aligned}$$

Using that $\sup_{h \in \mathcal{H}_1} P_n^1 \langle h, \phi \rangle_{\mathcal{H}}^2 \leq \|\phi\|_{L^2(P_n^1; \mathcal{H})}^2$ for $\phi : \mathcal{Z} \rightarrow \mathbb{H}$ and then subsequently applying the triangle inequality in $L^2(P_n^1; \mathcal{H})$, we find that

$$\begin{aligned}\|\Sigma_n^1 - \Sigma_0\|_{\text{op}} &\leq \|\phi_n^1 - \phi_0\|_{L^2(P_n^1; \mathcal{H})} (2\|\phi_0\|_{L^2(P_n^1; \mathcal{H})} + \|\phi_n^1 - \phi_0\|_{L^2(P_n^1; \mathcal{H})}) \\ &\quad + \sup_{h \in \mathcal{H}_1} (P_n^1 - P_0) \langle h, \phi_0 \rangle_{\mathcal{H}}^2.\end{aligned}$$

The second term is $o_p(1)$ by Lemma S13. We now show that the first term is also $o_p(1)$. To see this, first note that $\|\phi_0\|_{L^2(P_n^1; \mathcal{H})} = O_p(1)$ by the fact that $\|\phi_0\|_{L^2(P_0; \mathcal{H})} < \infty$ and since, by the weak law of large numbers, $\|\phi_0\|_{L^2(P_n^1; \mathcal{H})}^2 = \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 + o_p(1)$. Hence, it suffices to show that $\|\phi_n^1 - \phi_0\|_{L^2(P_n^1; \mathcal{H})} = o_p(1)$. To see that this holds, note that, for any $\delta > 0$, the probability that $\|\phi_n^1 - \phi_0\|_{L^2(P_n^1; \mathcal{H})}^2$ exceeds δ conditional on the data used to create the estimate ϕ_n^1 of ϕ_0 satisfies the following:

$$\begin{aligned}P_0^n \left\{ \|\phi_n^1 - \phi_0\|_{L^2(P_n^1; \mathcal{H})}^2 > \delta \mid Z_1, \dots, Z_{n/2} \right\} \\ \leq \min \left\{ 1, \delta^{-1} E_{P_0^n} \left[\|\phi_n^1 - \phi_0\|_{L^2(P_n^1; \mathcal{H})}^2 \mid Z_1, \dots, Z_{n/2} \right] \right\} \\ = \min \left\{ 1, \delta^{-1} \|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})}^2 \right\}.\end{aligned}$$

Taking an expectation of both sides over $Z_1, \dots, Z_{n/2}$, using that $\|\phi_n^1 - \phi_0\|_{L^2(P_0; \mathcal{H})}^2 = o_p(1)$, and applying the dominated convergence theorem shows that the right-hand side is $o(1)$. As $\delta > 0$ was arbitrary, this shows that $\|\phi_n^1 - \phi_0\|_{L^2(P_n^1; \mathcal{H})} = o_p(1)$, which gives the result. \square

E A conservative estimator of the threshold used to define our confidence sets that does not require the bootstrap

We now present a conservative estimator of the threshold $\zeta_{1-\alpha}$ that is used to construct the confidence sets described in Section 4.2. This estimator is applicable in settings where Ω_0 is the identity function. Its form is motivated by Theorem 1 in Székely and Bakirov (2003), which concerns tail probabilities for Gaussian quadratic forms of the type $\sum_{k=1}^{\infty} c_k N_k^2$, where $(N_k)_{k=1}^{\infty}$ is an iid sequence of standard normal random variables and $(c_k)_{k=1}^{\infty}$ is a sequence of nonnegative constants. This result is applicable when Ω_0 is the identity operator since, in that case, $\|\mathbb{H}\|_{\mathcal{H}}^2$ has the same distribution as $\sum_{k=1}^{\infty} E_0[\langle \phi_0(Z), h_{0,k} \rangle_{\mathcal{H}}^2] N_k^2$, where $(h_{0,k})_{k=1}^{\infty}$ are the unit eigenvectors of the covariance operator $E[\langle \mathbb{H}, \cdot \rangle_{\mathcal{H}} \mathbb{H}]$ of \mathbb{H} . When $\alpha \leq 0.2$, as it will be in most practical settings, Theorem 1 in Székely and Bakirov (2003) can be used to show that

$\Pr\{\|\mathbb{H}\|_{\mathcal{H}}^2 > \chi_{1-\alpha}^2 \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2\} \leq \alpha$, where $\chi_{1-\alpha}^2$ denotes the $(1-\alpha)$ -quantile of a chi-squared distribution with 1 degree of freedom. Hence, if s_n^2 is a consistent estimator of $\|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$, then Theorem 3 shows that $\mathcal{C}_n(\chi_{1-\alpha}^2 \cdot s_n^2)$ is an asymptotically valid, albeit conservative, $(1-\alpha)$ -confidence set for $\nu(P_0)$. If $\phi_n^j \rightarrow \phi_0$ in probability in $L^2(P_0; \mathcal{H})$ for $j \in \{1, 2\}$, then Lemma S14, given below, shows that the cross-fitted estimator $\frac{1}{2} \sum_{j=1}^2 \|\phi_n^j\|_{L^2(P_n^j; \mathcal{H})}^2$ will converge in probability to $\|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$, so that s_n^2 can be taken to be equal to this estimator.

Lemma S14. *If $\|\phi_0\|_{L^2(P_0; \mathcal{H})} < \infty$ and $\|\phi_n^j - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$ for each $j \in \{1, 2\}$, then $\|\phi_n^j\|_{L^2(P_n^j; \mathcal{H})}^2$ converges to $\|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$ in probability for each $j \in \{1, 2\}$ and, consequently, $\frac{1}{2} \sum_{j=1}^2 \|\phi_n^j\|_{L^2(P_n^j; \mathcal{H})}^2 \rightarrow \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$ in probability as well.*

Proof of Lemma S14. We begin by showing that $\|\phi_n^j\|_{L^2(P_n^j; \mathcal{H})}^2 \xrightarrow{p} \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$ for fixed $j \in \{1, 2\}$. Note that

$$\begin{aligned} & \left| \|\phi_n^j\|_{L^2(P_n^j; \mathcal{H})}^2 - \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 \right| \\ & \leq \left| \int \|\phi_n^j(z)\|_{\mathcal{H}}^2 (P_n^j - P_0)(dz) \right| + \left| \|\phi_n^j\|_{L^2(P_0; \mathcal{H})}^2 - \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 \right|. \end{aligned}$$

We study these two terms separately. The second term is $o_p(1)$ since (i) by the reverse triangle inequality and the assumption of this theorem, $|\|\phi_n^j\|_{L^2(P_0; \mathcal{H})} - \|\phi_0\|_{L^2(P_0; \mathcal{H})}| \leq \|\phi_n^j - \phi_0\|_{L^2(P_0; \mathcal{H})} = o_p(1)$ and (ii) by the continuous mapping theorem, $\|\phi_n^j\|_{L^2(P_0; \mathcal{H})} \xrightarrow{p} \|\phi_0\|_{L^2(P_0; \mathcal{H})}$ implies that $\|\phi_n^j\|_{L^2(P_0; \mathcal{H})}^2 \xrightarrow{p} \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$. In what follows, we show that the first term above is $o_p(1)$ as well.

Combining the fact that $\|\phi_n^j(z)\|_{\mathcal{H}} \leq \|\phi_n^j(z) - \phi_0(z)\|_{\mathcal{H}} + \|\phi_0(z)\|_{\mathcal{H}}$ with the basic inequality that $(a+b)^2 \leq 2(a^2 + b^2)$, and subsequently applying the triangle inequality, yields that

$$\begin{aligned} & \frac{1}{2} \left| \int \|\phi_n^j(z)\|_{\mathcal{H}}^2 (P_n^j - P_0)(dz) \right| \\ & \leq \left| \int (\|\phi_n^j(z) - \phi_0(z)\|_{\mathcal{H}}^2 + \|\phi_0(z)\|_{\mathcal{H}}^2) (P_n^j - P_0)(dz) \right| \\ & \leq \left| \int \|\phi_n^j(z) - \phi_0(z)\|_{\mathcal{H}}^2 P_n^j(dz) \right| + \left| \int \|\phi_n^j(z) - \phi_0(z)\|_{\mathcal{H}}^2 P_0(dz) \right| \\ & \quad + \left| \int \|\phi_0(z)\|_{\mathcal{H}}^2 (P_n^j - P_0)(dz) \right|. \end{aligned} \tag{S35}$$

The second term is equal to $\|\phi_n^j(z) - \phi_0(z)\|_{L^2(P_0; \mathcal{H})}^2$ and so is $o_p(1)$ by assumption. The third term is $o_p(1)$ by the weak law of large numbers, which is applicable since $\|\phi_0\|_{L^2(P_0; \mathcal{H})}^2 < \infty$ by assumption. By Markov's inequality, the fact that P_n^j and \hat{P}_n^j are fitted on different subsamples, and the fact that probabilities are no more than 1, the conditional probability that the first term exceeds any fixed $\delta > 0$ satisfies the following:

$$\begin{aligned} & P_0^n \left\{ \left| \int \|\phi_n^j(z) - \phi_0(z)\|_{\mathcal{H}}^2 P_n^j(dz) \right| > \delta \middle| Z_{(j-1)n/2+1}, \dots, Z_{(j-1)n/2+n/2} \right\} \\ & \leq \min \left\{ 1, \frac{1}{\delta} \|\phi_n^j - \phi_0\|_{L^2(P_0; \mathcal{H})}^2 \right\}. \end{aligned}$$

Taking an expected value of both sides over $Z_{(j-1)n/2+1}, \dots, Z_{(j-1)n/2+n/2} \stackrel{\text{iid}}{\sim} P_0$ and using that $\|\phi_n^j(z) - \phi_0(z)\|_{L^2(P_0; \mathcal{H})}^2 = o_p(1)$, the dominated convergence theorem shows that $\left| \int \|\phi_n^j(z) - \phi_0(z)\|_{\mathcal{H}}^2 P_n^j(dz) \right| > \delta$ occurs with probability tending to zero. As $\delta > 0$ was arbitrary, this shows that the first term on the right-hand side of (S35) is $o_p(1)$, which completes the proof of the fact that $\|\phi_n^j\|_{L^2(P_n^j; \mathcal{H})}^2 \xrightarrow{p} \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$ for $j \in \{1, 2\}$.

Since $\|\phi_n^j\|_{L^2(P_n^j; \mathcal{H})}^2 \xrightarrow{p} \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$ for $j \in \{1, 2\}$, the continuous mapping theorem shows that $\frac{1}{2} \sum_{j=1}^2 \|\phi_n^j\|_{L^2(P_n^j; \mathcal{H})}^2 \rightarrow \|\phi_0\|_{L^2(P_0; \mathcal{H})}^2$. \square

F Numerical considerations for computing the proposed confidence sets

Evaluating whether some $h_0 \in \mathcal{H}$ belongs to the confidence set in (24) requires computing the quadratic form $\langle \Omega_n(\bar{\nu}_n - h_0), \bar{\nu}_n - h_0 \rangle_{\mathcal{H}}$ in a possibly infinite dimensional Hilbert space. In many cases, computing this quadratic form will require leveraging some form of numerical approximation. One way of doing this is to replace the computation of the quadratic form in (24) by a finite-dimensional approximation thereof. To this end, for each $m \in \mathbb{N}$ we let $D_m : \mathcal{H} \rightarrow \mathbb{R}^m$ denote a linear operator. This linear operator should have the property that, for any h_1 and h_2 in \mathcal{H} , $\langle h_1, h_2 \rangle_{\mathcal{H}} \xrightarrow{m \rightarrow \infty} D_m(h_1)^\top D_m(h_2)$, where here and in all subsequent calculations all vectors are taken to be equal to column vectors when involved in matrix operations. In practice, for a given sample size n , m can be chosen to be some large constant. One natural choice of D_m corresponds to the map from h to the vector of the first m generalized Fourier coefficients of h with respect to some orthonormal basis $(h_k)_{k=1}^\infty$, so that $D_m(h) = (\langle h_k, h \rangle_{\mathcal{H}})_{k=1}^m$. If $\mathcal{H} = L^2([0, 1])$ and it is known that D_m will only be evaluated on elements of \mathcal{H} that have a continuous version, as occurs if h_0 is continuous and $\bar{\nu}_n$ is continuous with probability one, then another natural choice is to take $D_m(h)$ to be equal to $(h(t_k)/m^{1/2})_{k=1}^m$, where $t_k = k/(m+1)$ and $h(t_k)$ is taken to be the evaluation of the continuous version of h at t_k . If instead $\mathcal{H} = L^2(\mathbb{R})$ and h still has a continuous version, then D_m can be taken equal to $(h(t_k)/[m\varphi_{\mu, \sigma}(t_k)]^{1/2})_{k=1}^m$, where, for $\mu \in \mathbb{R}$ and $\sigma > 0$, $\{t_k\}_{k=1}^m$ are such that $\Phi_{\mu, \sigma}(t_k) = k/(m+1)$ with $\Phi_{\mu, \sigma}$ and $\varphi_{\mu, \sigma}$ denoting the cumulative distribution function and probability density function of a $N(\mu, \sigma^2)$ distribution, respectively. In practice μ and σ may be selected based on the data, which can be justified theoretically so long as their random values converge to some limits in probability asymptotically — for example, in our simulation implementation of the bandlimited density estimator from Example 1b, we take μ and $\sigma/4$ to be the empirical mean and standard deviation of Y given $A = 1$, respectively.

The linear operator D_m can be used to approximate the infinite-dimensional quadratic form in (24) with a finite-dimensional quadratic form. In particular, $\langle \Omega_n(\bar{\nu}_n - h_0), \bar{\nu}_n - h_0 \rangle_{\mathcal{H}}$ can be replaced by $D_m(\bar{\nu}_n - h_0)^\top \tilde{\Omega}_{n,m} D_m(\bar{\nu}_n - h_0)$, where $\tilde{\Omega}_{n,m}$ is an m -dimensional positive definite Hermitian matrix whose value will depend on the standardization operator Ω_n that it is meant to approximate. If Ω_n is the identity operator, then $\tilde{\Omega}_{n,m}$ can be taken to be equal to the m -dimensional identity matrix I_m . If Ω_n is instead the estimator of the regularized covariance operator described in Appendix D, then it can instead be approximated by a regularized empirical covariance matrix. In particular, we can let $\tilde{\Omega}_{n,m} := [(1 - \lambda)\Sigma_{n,m} + \lambda I_m]^{-1}$, where $\Sigma_{n,m} := \frac{1}{2} \sum_{j=1}^2 P_n^j [D_m(\phi_n^j(\cdot)) D_m(\phi_n^j(\cdot))^\top]$, where $P_n^j [D_m(\phi_n^j(\cdot)) D_m(\phi_n^j(\cdot))^\top]$ corresponds to the empirical covariance matrix of the random variable $D_m(\phi_n^j(Z))$ computed using the empirical distribution P_n^j . Alternatively, $P_n^j [D_m(\phi_n^j(\cdot)) D_m(\phi_n^j(\cdot))^\top]$ may be replaced by the empirical correlation matrix of $D_m(\phi_n^j(Z))$ under P_n^j in the definition of $\tilde{\Omega}_{n,m}$. Though using an empirical correlation matrix rather than an empirical covariance matrix changes the quadratic form used to define the confidence set, doing so can make selecting the parameter λ simpler because, in that case, the matrices $\Sigma_{n,m}$ and I_m are on the same scale in the sense that both have trace m .

We conclude by noting that, when \mathcal{H} is an RKHS on \mathcal{T} with feature map $t \mapsto K_t$, it will be possible to compute the quadratic form $\langle \Omega_n(\bar{\nu}_n - h_0), \bar{\nu}_n - h_0 \rangle_{\mathcal{H}}$ explicitly in some cases. One particularly interesting case occurs when $\bar{\nu}_n$ and h_0 are both contained in the linear span of $\{K_{t_k}\}_{k=1}^m$ and Ω_n is the identity operator, where the set $\{K_{t_k}\}_{k=1}^m$ may depend on the observed data. In such cases, $\bar{\nu}_n - h_0 = \sum_{k=1}^m c_k K_{t_k}$ for some $c := (c_k)_{k=1}^m \in \mathbb{R}^m$, and so, letting G denote the Gram matrix with $G_{jk} = K_{t_j}(t_k)$, it holds that

$\langle \bar{\nu}_n - h_0, \bar{\nu}_n - h_0 \rangle_{\mathcal{H}} = c^\top G c$. If $h_0 = 0$, which would be the key value of h_0 to consider when the confidence set is being used to test the null hypothesis that $\nu(P_0) = 0$ against the complementary alternative, it is necessarily the case that h_0 is in the linear span of $\{K_{t_k}\}_{k=1}^m$ for any collection $\{t_k\}_{k=1}^m$. Hence, in these cases, it suffices that $\bar{\nu}_n$ be in the linear span of $\{K_{t_k}\}_{k=1}^m$ for some m . Such cases arise, for example, when using the MMD to test for the equality of two distributions (Gretton et al., 2012); when conducting these tests, $\bar{\nu}_n$ is equal to the difference of the one-step estimators of the kernel mean embeddings of two distributions.

G Cross-validated selection of the regularization parameter

The key observation that motivates the risk we use is that, along any quadratic mean differentiable submodel $\{P_\epsilon : \epsilon \in [0, \delta]\} \in \mathcal{P}(P, \mathcal{P}, s)$, $\nu(P_\epsilon) - \nu(P)$ should approximately be equal to $\epsilon \dot{\nu}_P(s)$ or, put another way, should approximately be equal to

$$\begin{aligned} \operatorname{argmin}_{h \in \mathcal{H}} \|h - \epsilon \dot{\nu}_P(s)\|_{\mathcal{H}}^2 &= \operatorname{argmin}_{h \in \mathcal{H}} \left[\frac{1}{2} \|h\|_{\mathcal{H}}^2 - \epsilon \langle h, \dot{\nu}_P(s) \rangle_{\mathcal{H}} \right] \\ &= \operatorname{argmin}_{h \in \mathcal{H}} \left[\frac{1}{2} \|h\|_{\mathcal{H}}^2 - \epsilon \langle \dot{\nu}_P^*(h), s \rangle_{L^2(P)} \right]. \end{aligned}$$

By Lemma S6 in Appendix C, $\epsilon \langle \dot{\nu}_P^*(h), s \rangle_{L^2(P)} = P_\epsilon \dot{\nu}_P^*(h) + o(\epsilon)$ under appropriate conditions. This suggests that $\nu(P_\epsilon) - \nu(P)$ should approximately equal $\operatorname{argmin}_{h \in \mathcal{H}} [\frac{1}{2} \|h\|_{\mathcal{H}}^2 - P_\epsilon \dot{\nu}_P^*(h)]$. Letting P_0 play the role of P_ϵ and an estimate \hat{P}_n^{loss} of P_0 play the role of P suggests that $\nu(P_0) - \nu(\hat{P}_n^{\text{loss}})$ should approximately minimize $\frac{1}{2} \|h\|_{\mathcal{H}}^2 - P_0 \dot{\nu}_{\hat{P}_n^{\text{loss}}}(h)$ over $h \in \mathcal{H}$. Put another way, $\nu(P_0)$ should approximately minimize $E_0[\mathcal{L}_{\hat{P}_n^{\text{loss}}}(Z; h)] = \frac{1}{2} \|h - \nu(\hat{P}_n^{\text{loss}})\|_{\mathcal{H}}^2 - P_0 \dot{\nu}_{\hat{P}_n^{\text{loss}}}[h - \nu(\hat{P}_n^{\text{loss}})]$ over $h \in \mathcal{H}$. This suggests using the loss $\mathcal{L}_{\hat{P}_n^{\text{loss}}}$ when performing cross-validation to select the regularization parameter β_n . Such an approach is presented in Algorithm 1. We refer the reader to van der Laan and Dudoit (2003) for arguments that can be used to establish oracle guarantees for this cross-validation selector.

H Inefficient influence operators and influence functions

When \mathcal{H} is finite-dimensional and the model is semiparametric at P_0 — in the sense that its tangent space $\dot{\mathcal{P}}_{P_0}$ is a strict subspace of $L_0^2(P_0)$ — there are generally many influence functions that can be used to construct a one-step estimator of $\nu(P_0)$. In our more general Hilbert-valued setting, the same can be done by replacing the efficient influence operator that we use to construct our (regularized) one-step estimators by an inefficient influence operator. Each inefficient influence operator is the Hermitian adjoint of a bounded linear extension $\dot{\nu}_{\text{ext}, P}$ of the local parameter $\dot{\nu}_P$ from $\dot{\mathcal{P}}_P$ to $L_0^2(P)$. All such extensions take the form $\dot{\nu}_{\text{ext}, P}(s) = \dot{\nu}_P(\Pi_{L_0^2(P)}[s \mid \dot{\mathcal{P}}_P]) + \dot{\xi}_P(\Pi_{L_0^2(P)}[s \mid \dot{\mathcal{P}}_P^\perp])$, where $\dot{\mathcal{P}}_P^\perp$ is the orthogonal complement of $\dot{\mathcal{P}}_P \subset L_0^2(P)$ and $\dot{\xi}_P : \dot{\mathcal{P}}_P^\perp \rightarrow \mathcal{H}$ is bounded and linear. The corresponding influence operator is the Hermitian adjoint of $\dot{\nu}_{\text{ext}, P}$, which takes the form $\dot{\nu}_{\text{ext}, P}^* = \dot{\nu}_P^* + \dot{\xi}_P^*$, where $\dot{\nu}_P^*$ is the efficient influence operator and $\dot{\xi}_P^*$ is the Hermitian adjoint of $\dot{\xi}_P$. The efficient influence operator is recovered by taking $\dot{\xi}_P$ to be the zero operator. If $\dot{\xi}_P \neq 0$, which we assume hereafter, then $\dot{\nu}_{\text{ext}, P}^* \neq \dot{\nu}_P^*$ and we call $\dot{\nu}_{\text{ext}, P}^*$ an inefficient influence operator.

In some cases, $\dot{\nu}_{\text{ext}, P}^*$ will have an associated inefficient influence function $\phi_{\text{ext}, P}$. Concretely, this holds if and only if $\dot{\nu}_{\text{ext}, P}^*(\cdot)(z)$ is bounded and linear P -almost surely; in these cases, $\phi_{\text{ext}, P}(z)$ is the Riesz representation of this operator. The following result shows that an efficient influence function must exist for an inefficient one to exist, and also provides a means to derive the form of the EIF based on the form of an inefficient influence function. We let Φ denote the $L^2(P; \mathcal{H})$ -closure of the linear span of

Algorithm 1 Cross-validated selection of the regularization parameter β_n

- 1: **Inputs:** Data Z_1, Z_2, \dots, Z_n , estimator to be used to estimate the nuisance P_0 and a finite subset B_n of ℓ_*^2 of candidate values for the regularization parameter
- 2: **Generate folds:** partition the multiset $\{Z_i\}_{i=1}^n$ into multisets $\mathcal{Z}_1, \mathcal{Z}_2, \mathcal{Z}_3, \mathcal{Z}_4$ of roughly equal size
- 3: **for** all folds $j = 1, 2, 3, 4$ **do**
- 4: **Nuisance estimation:** using only data in \mathcal{Z}_j , estimate P_0 as \hat{P}_n^j .
- 5: **end for**
- 6: **for** all permutations $j = (j(1), j(2), j(3), j(4))$ of $\{1, 2, 3, 4\}$ **do**
- 7: **Nuisance for regularized one-step:** let $\hat{P}_n^{\text{os}} = \hat{P}_n^{j(1)}$
- 8: **Nuisance for loss function:** let $\hat{P}_n^{\text{loss}} = \hat{P}_n^{j(2)}$
- 9: **for** all candidate regularization parameters $\beta := (\beta_k)_{k=1}^\infty \in B_n$ **do**
- 10: **Define regularized one-step estimator:** $\tilde{\nu}_{n,j}^\beta := \nu(\hat{P}_n^{\text{os}}) + \frac{1}{|\mathcal{Z}_{j(3)}|} \sum_{z \in \mathcal{Z}_{j(3)}} \phi_{\hat{P}_n^{\text{os}}}^\beta(z)$
- 11: **Compute risk** (j, β) -specific risk: $R_j^\beta := \frac{1}{|\mathcal{Z}_{j(4)}|} \sum_{z \in \mathcal{Z}_{j(4)}} \mathcal{L}_{\hat{P}_n^{\text{loss}}}(z; \tilde{\nu}_{n,j}^\beta)$
- 12: **end for**
- 13: **end for**
- 14: **for** all candidate regularization parameters $\beta := (\beta_k)_{k=1}^\infty \in B_n$ **do**
- 15: **Aggregate the risks:** $R^\beta := \frac{1}{24} \sum_{\text{permutations } j \text{ of } \{1, 2, 3, 4\}} R_j^\beta$
- 16: **end for**
- 17: **return** $\beta^* \in \operatorname{argmin}_{\beta \in B_n} R^\beta$

$$\{z \mapsto s(z)h : h \in \mathcal{H}, s \in \dot{\mathcal{P}}_P\}.$$

Lemma S15 (Expressing the EIF in terms of an inefficient influence function). *If ν is pathwise differentiable at $P \in \mathcal{P}$ with inefficient influence function $\phi_{\text{ext},P} \in L^2(P; \mathcal{H})$, then ν has EIF $\phi_P = \Pi_{L^2(P; \mathcal{H})} [\phi_{\text{ext},P} \mid \Phi] \in L^2(P; \mathcal{H})$.*

The proof of this lemma is given at the end of this appendix.

One-step estimators can be constructed using inefficient influence operators. Beginning with cases where an inefficient influence function exists and using notation from Section 2.5, we define a cross-fitted one-step estimator as $\bar{\nu}_{\text{ext},n} := \frac{1}{2} \sum_{j=1}^2 [\nu(\hat{P}_n^j) + P_n^j \phi_{\text{ext},n}^j]$, where $\phi_{\text{ext},n}^j := \phi_{\text{ext},\hat{P}_n^j}$. Under similar conditions to those of Theorem 2, it can be shown that $\bar{\nu}_{\text{ext},n}$ is regular and asymptotically linear with influence function $\phi_{\text{ext},0} := \phi_{\text{ext},P_0}$ and

$$n^{1/2} [\bar{\nu}_{\text{ext},n} - \nu(P_0)] \rightsquigarrow \mathbb{H}_{\text{ext}}, \quad (\text{S36})$$

where \mathbb{H}_{ext} is a tight \mathcal{H} -valued Gaussian random variable that is such that, for each $h \in \mathcal{H}$, $\langle \mathbb{H}_{\text{ext}}, h \rangle_{\mathcal{H}} \sim N(0, E_0[\langle \phi_{\text{ext},0}(Z), h \rangle_{\mathcal{H}}^2])$. The above weak convergence facilitates the construction of confidence sets for $\nu(P_0)$ using analogous methods to those used in Section 4.2. The main distinction between $\bar{\nu}_{\text{ext},n}$ and $\bar{\nu}_n$ is that, since $\phi_{\text{ext},0}$ is not the EIF, the conditions of the convolution theorem fail to hold (Theorem 3.11.2 and Lemma 3.11.4 of van der Vaart and Wellner, 1996), and so $\bar{\nu}_{\text{ext},n}$ will not be efficient — e.g., (23) will not generally hold.

Moving now to cases where an inefficient influence function does not exist, we define a β_n -regularized one-step estimator as $\bar{\nu}_{\text{ext},n}^{\beta_n} := \frac{1}{2} \sum_{j=1}^2 [\nu(\hat{P}_n^j) + P_n^j \phi_{\text{ext},n}^{j,\beta_n}]$ with $\phi_{\text{ext},n}^{j,\beta_n}(z) := \sum_{k=1}^\infty \beta_{n,k} \dot{\nu}_{\text{ext},\hat{P}_n^j}^*(h_k)(z) h_k$. This regularized one-step estimator satisfies similar guarantees to those satisfied by the one based on

the efficient influence operator: it achieves a $\|\beta_n\|_{\ell^2}/n^{1/2}$ -rate of convergence when a drift, regularized remainder, and bias terms are small; the drift term will be small if $\phi_{\text{ext},n}^{j,\beta_n}$ is close to $\phi_{\text{ext},0}^{\beta_n}(z) := \sum_{k=1}^{\infty} \beta_{n,k} \dot{\nu}_{\text{ext},P_0}^*(h_k)(z) h_k$ in $L^2(P_0; \mathcal{H})$; the remainder will be small if $\max_j \sup_{k \in \mathbb{N}} |\langle \nu(\widehat{P}_n^j) - \nu(P_0), h_k \rangle_{\mathcal{H}} + P_0 \dot{\nu}_{\text{ext},\widehat{P}_n^j}^*(h_k)|$ is $o_p(n^{-1/2})$; and the bias term will be small if $\nu(P)$ is sufficiently smooth for all $P \in \mathcal{P}$. As for confidence sets, the same methods as described in Section 5.2 can be used once one notes that, for any fixed $\beta \in \ell^2 \cap (0, 1]^{\mathbb{N}}$, $\nu^{\beta} := \Gamma_{\beta} \circ \nu$ is pathwise differentiable at P_0 with inefficient influence function $\phi_{\text{ext},0}^{\beta}$. This can be used to justify, for example, the asymptotic validity of the $(1 - \alpha)$ confidence set

$$\left\{ h \in \mathcal{H} : \sum_{k=1}^{\infty} \beta_k^2 \left[\frac{1}{2} \sum_{j=1}^2 \left\{ \langle \nu(\widehat{P}_n^j), h_k \rangle_{\mathcal{H}} + P_n^j \dot{\nu}_{\text{ext},\widehat{P}_n^j}^*(h_k) \right\} - \langle h, h_k \rangle_{\mathcal{H}} \right]^2 \leq \widehat{\zeta}_{\text{ext},n}/n \right\},$$

where $\widehat{\zeta}_{\text{ext},n}$ is selected via the bootstrap. Since $\dot{\nu}_{\text{ext},\widehat{P}_n^j}^*$ is an inefficient influence operator, the threshold $\widehat{\zeta}_{\text{ext},n}$ will generally be asymptotically larger than the one used for the confidence set built based on the efficient influence operator given in (30) (see Lemma 3.11.4 of [van der Vaart and Wellner, 1996](#)).

Proof of Lemma S15. Denote by $\dot{\nu}_{\text{ext},P}^*$ the inefficient influence operator to which $\phi_{\text{ext},P}$ corresponds, and let its Hermitian adjoint $\dot{\nu}_{\text{ext},P}$ denote the extension of the local parameter $\dot{\nu}_P$ used to define this inefficient influence operator. Throughout this proof we let $\phi_P^{\diamond} := \Pi_{L^2(P; \mathcal{H})}[\phi_{\text{ext},P} \mid \Phi] \in L^2(P; \mathcal{H})$. Our goal is to show that ν has EIF $\phi_P = \phi_P^{\diamond}$.

Since we have assumed throughout that a separable version of the efficient influence process is used, there exists a countable dense subset \mathcal{H}' of \mathcal{H} and a P -probability one subset \mathcal{Z}' of \mathcal{Z} such that, for all $h \in \mathcal{H}$ and $z \in \mathcal{Z}'$, there exists an \mathcal{H}' -valued sequence $(h'_j)_{j=1}^{\infty}$ that converges to h and satisfies $\dot{\nu}_P^*(h'_j)(z) \rightarrow \dot{\nu}_P^*(h)(z)$ as $j \rightarrow \infty$. Fix $\epsilon > 0$ and $z \in \mathcal{Z}'$. Let h_{ϵ} be such that

$$|\dot{\nu}_P^*(h_{\epsilon})(z) - \langle h_{\epsilon}, \phi_P^{\diamond}(z) \rangle_{\mathcal{H}}| \geq \sup_{h \in \mathcal{H}} |\dot{\nu}_P^*(h)(z) - \langle h, \phi_P^{\diamond}(z) \rangle_{\mathcal{H}}| - \epsilon.$$

Fix an \mathcal{H}' -valued sequence $(h_{\epsilon,j})_{j=1}^{\infty}$ that converges to h_{ϵ} and is such that $\dot{\nu}_P^*(h_{\epsilon,j})(z) \rightarrow \dot{\nu}_P^*(h_{\epsilon})(z)$ as $j \rightarrow \infty$. By the choice of this sequence and the fact that $\langle \cdot, \phi_P^{\diamond}(z) \rangle_{\mathcal{H}}$ is continuous, there exists a large enough j such that

$$|\dot{\nu}_P^*(h_{\epsilon,j})(z) - \langle h_{\epsilon,j}, \phi_P^{\diamond}(z) \rangle_{\mathcal{H}}| \geq \sup_{h \in \mathcal{H}} |\dot{\nu}_P^*(h)(z) - \langle h, \phi_P^{\diamond}(z) \rangle_{\mathcal{H}}| - 2\epsilon.$$

Taking a supremum over $h_{\epsilon,j} \in \mathcal{H}'$ on the left and then recalling that $\epsilon > 0$ was arbitrary shows that

$$\sup_{h \in \mathcal{H}'} |\dot{\nu}_P^*(h)(z) - \langle h, \phi_P^{\diamond}(z) \rangle_{\mathcal{H}}| = \sup_{h \in \mathcal{H}} |\dot{\nu}_P^*(h)(z) - \langle h, \phi_P^{\diamond}(z) \rangle_{\mathcal{H}}|.$$

Now, for each $h \in \mathcal{H}'$, let $\mathcal{Z}_h'' := \{z \in \mathcal{Z} : \dot{\nu}_P^*(h)(z) = \langle h, \phi_P^{\diamond}(z) \rangle_{\mathcal{H}}\}$. In the remainder we shall show that \mathcal{Z}_h'' has P -probability one. Combining this with the above display then shows that ν has EIF $\phi_P = \phi_P^{\diamond}$, since that will show that the right-hand side above is 0 on the P -probability one set $\mathcal{Z}' \cap [\cap_{h \in \mathcal{H}'} \mathcal{Z}_h'']$.

Fix $h \in \mathcal{H}'$ and let $\langle h, \phi_P^{\diamond} \rangle_{\mathcal{H}}$ denote the function $z \mapsto \langle h, \phi_P^{\diamond}(z) \rangle_{\mathcal{H}}$. We will show that $\langle \dot{\nu}_P^*(h) - \langle h, \phi_P^{\diamond} \rangle_{\mathcal{H}}, u \rangle_{L^2(P)} = 0$ for a generic $u \in L^2(P)$, which will then establish that \mathcal{Z}_h'' is a P -probability one set and complete our proof. Writing $u = s + s^{\perp}$ with $s \in \dot{\mathcal{P}}_P$ and s^{\perp} belonging to the orthogonal complement of $\dot{\mathcal{P}}_P$ in $L^2(P)$, it suffices to show that $\langle \dot{\nu}_P^*(h) - \langle h, \phi_P^{\diamond} \rangle_{\mathcal{H}}, s \rangle_{L^2(P)} = 0$ and $\langle \dot{\nu}_P^*(h) - \langle h, \phi_P^{\diamond} \rangle_{\mathcal{H}}, s^{\perp} \rangle_{L^2(P)} = 0$. Beginning with the former equality and using that the restriction of $\dot{\nu}_{\text{ext},P}$ to $\dot{\mathcal{P}}_P$ is equal to $\dot{\nu}_P$,

$$\begin{aligned} \langle \dot{\nu}_P^*(h), s \rangle_{L^2(P)} &= \langle h, \dot{\nu}_P(s) \rangle_{\mathcal{H}} = \langle h, \dot{\nu}_{\text{ext},P}(s) \rangle_{\mathcal{H}} = \langle \dot{\nu}_{\text{ext},P}^*(h), s \rangle_{L^2(P)} \\ &= \int \dot{\nu}_{\text{ext},P}^*(h)(z) s(z) P(dz) = \int \langle h, \phi_{\text{ext},P}(z) \rangle_{\mathcal{H}} s(z) P(dz) \end{aligned}$$

$$= \int \langle s(z)h, \phi_{\text{ext},P}(z) \rangle_{\mathcal{H}} P(dz) = \langle z \mapsto s(z)h, \phi_{\text{ext},P} \rangle_{L^2(P; \mathcal{H})}.$$

Since $z \mapsto s(z)h \in \Phi$ and ϕ_P^\diamond is an orthogonal projection of $\phi_{\text{ext},P}$ onto Φ , the right-hand side equals $\langle z \mapsto s(z)h, \phi_P^\diamond \rangle_{L^2(P; \mathcal{H})}$. By similar calculations to those used above, this in turn equals $\langle \langle h, \phi_P^\diamond \rangle_{\mathcal{H}}, s \rangle_{L^2(P)}$. Hence, $\langle \dot{\nu}_P^*(h) - \langle h, \phi_P^\diamond \rangle_{\mathcal{H}}, s \rangle_{L^2(P)} = 0$. To see that $\langle \dot{\nu}_P^*(h) - \langle h, \phi_P^\diamond \rangle_{\mathcal{H}}, s^\perp \rangle_{L^2(P)} = 0$, observe that $\langle \dot{\nu}_P^*(h), s^\perp \rangle_{L^2(P)} = 0$ since $\dot{\nu}_P^*$ has codomain $\dot{\mathcal{P}}_P$, and

$$\langle \langle h, \phi_P^\diamond \rangle_{\mathcal{H}}, s^\perp \rangle_{L^2(P)} = \langle z \mapsto s^\perp(z)h, \phi_P^\diamond \rangle_{L^2(P; \mathcal{H})} = 0,$$

where we have used that all elements of $\Phi \subset L^2(P; \mathcal{H})$ are orthogonal to $z \mapsto s^\perp(z)h$ by virtue of the fact that s^\perp is orthogonal to $\dot{\mathcal{P}}_P \subset L^2(P)$. \square

I Additional simulation results

I.1 Simulation results for Example 1b

We evaluate the coverage of our spherical $L^2(\mathbb{R})$ confidence sets for a bandlimited counterfactual density of $Y(1)$. When doing this, we take $Q(1)$ to be the distribution of $\sigma_M S + \mu_M$, where $(\mu_1, \mu_2, \mu_3) = (-6, -2, 5)$, $(\sigma_1, \sigma_2, \sigma_3) = (2, 3, 2)$, and $M \sim \text{Unif}\{1, 2, 3\}$ is drawn independently of the random variable S , which has density function $3\text{sinc}^4(\cdot)/(2\pi)$. The density of $Q(1)$ is depicted in the top left corner of Figure 1. It is bandlimited, with the support of its Fourier transform equal to $[-2, 2]$. We focus on the case where the bandlimiting parameter b used to define $\underline{\nu}$ in Eq. 15 is equal to 2, so that the target of inference $\underline{\nu}(P_0)$ coincides with the density of $Q(1)$.

Figure S1 displays the coverage of our confidence sets at different nominal levels. At all nominal levels larger than 75%, the confidence sets are slightly conservative at the sample sizes considered, with the actual coverage probability approaching the nominal level as n grows. A similar improvement with n holds across the full 0-100% range of nominal levels, which both supports our theoretical weak convergence guarantees for the one-step estimator and their finite-sample utility. In Table S1, we also verified that, as anticipated by our theory for cases where an EIF exists, the mean integrated squared error of the one-step estimator decays at an n^{-1} rate.

We conclude by comparing the size of our $L^2(\mathbb{R})$ confidence set \mathcal{C}_n to those of a pointwise confidence interval for the counterfactual density function at zero, namely $\underline{\nu}(P_0)(0)$. To make this comparison, we first note that, for any $y \in \mathbb{R}$, including $y = 0$, and any $h \in \mathcal{C}_n := \{h \in \mathcal{H} : \|h - \bar{\nu}_n\|_{L^2(\lambda_Y)}^2 \leq \hat{\zeta}_n/n\}$,

$$\begin{aligned} |h(y) - \bar{\nu}_n(y)| &= \left| \int_{-\infty}^{\infty} K_y(\tilde{y}) [h(\tilde{y}) - \bar{\nu}_n(\tilde{y})] \lambda_Y(d\tilde{y}) \right| \\ &\leq \|K_y\|_{L^2(\lambda_Y)} \|h - \bar{\nu}_n\|_{L^2(\lambda_Y)} \leq (b\hat{\zeta}_n/[n\pi])^{1/2}, \end{aligned} \quad (\text{S37})$$

where we used that $\|K_y\|_{L^2(\lambda_Y)} = (b/\pi)^{1/2}$. Consequently, our $L^2(\mathbb{R})$ confidence set yields an interval for $\underline{\nu}(P_0)(0)$ of the form $\bar{\nu}_n(0) \pm (b\hat{\zeta}_n/[n\pi])^{1/2}$. In our simulation setting, this confidence interval was about 2.3 times wider than an efficient Wald-type confidence interval for the real-valued quantity $\underline{\nu}(P_0)(0)$ when $\alpha = 0.05$. Hence, if a point evaluation of $\underline{\nu}(P_0)$ is truly the target of inference, then there would be a benefit to directly pursuing inference for this quantity, rather than the function as a whole. However, if the function $\underline{\nu}(P_0)$ is the target of inference, then our $L^2(\mathbb{R})$ confidence set is likely the preferred method for making inference. This may be especially true in this bandlimited density example since the fact that (S37) holds for all $y \in \mathbb{R}$ shows that a uniform confidence band for $\underline{\nu}(P_0)$ is given by $\mathcal{C}_{n,\infty} := \{h \in \mathcal{H} : \sup_{y \in \mathbb{R}} |h(y) - \bar{\nu}_n(y)| \leq (b\hat{\zeta}_n/[n\pi])^{1/2}\}$.

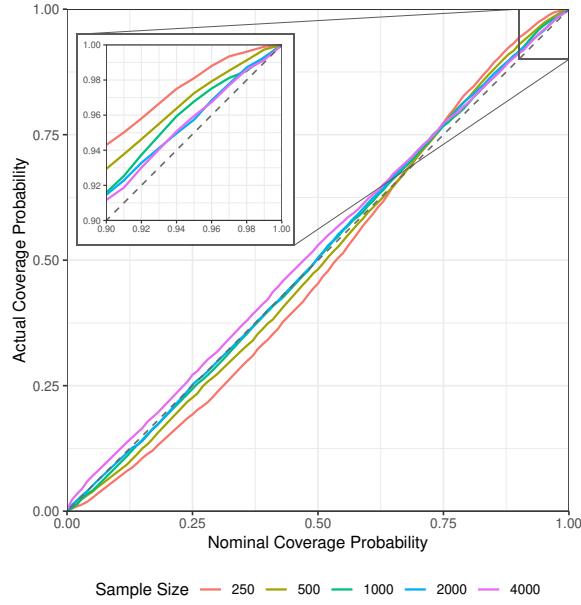


Figure S1: Actual versus nominal coverage of confidence sets for the bandlimited density function based on 5000 Monte Carlo repetitions. The inset displays nominal coverage values that are of particular interest in practice.

	Sample Size (n)				
	250	500	1000	2000	4000
Plug-In	1.11	1.59	2.22	3.04	4.15
One-Step	3.58	3.30	3.11	2.97	2.84

Table S1: Performance of the plug-in and one-step estimators at different sample sizes n , where performance is measured in terms of n times the mean integrated squared error. As would be predicted by theory, this criterion appears to stabilize with n for the one-step estimator. In contrast, it grows with n for the plug-in estimator.

I.2 Supplemental tables and figures for simulation studies

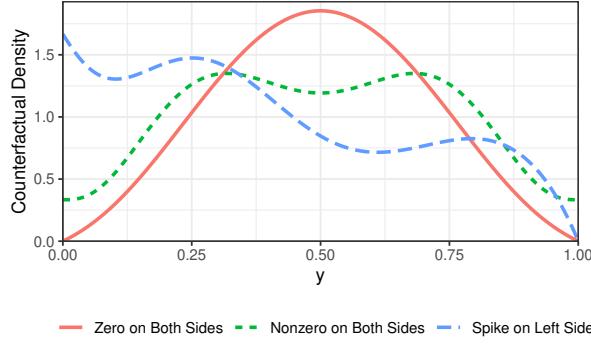
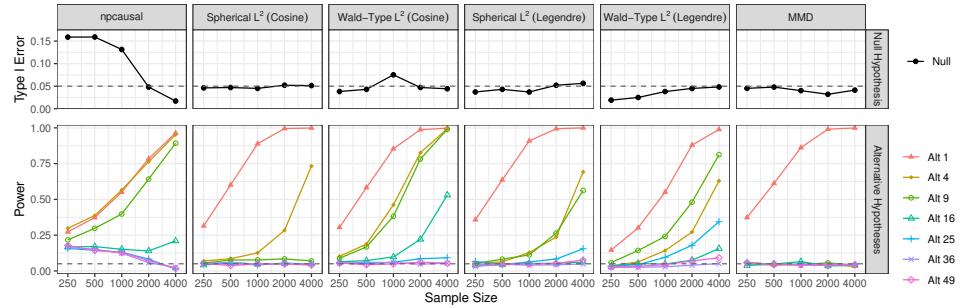
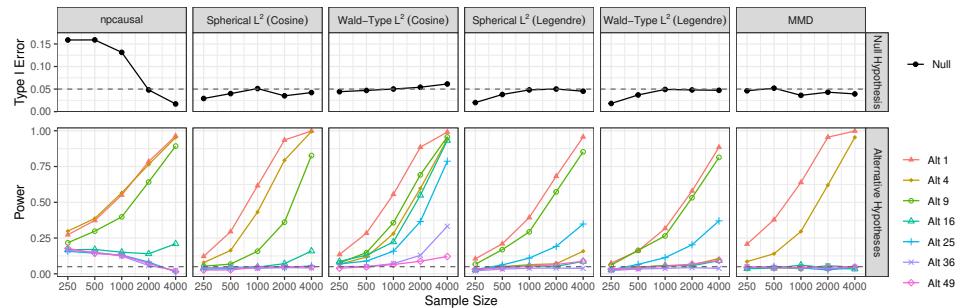


Figure S2: Densities $Q(1)$ of $Y(1)$ used in the three settings considered in the evaluation of estimators in Example 1a. Each of these densities is a uniform mixture $\frac{1}{3} \sum_{k=1}^3 \text{Beta}(c_k, d_k)$, where the parameters indexing the beta distributions in this mixture differ across the three settings plotted in the figure. In particular, $\{(c_1, d_1), (c_2, d_2), (c_3, d_3)\}$ is equal to $\{(2, 2), (3, 3), (4, 4)\}$ for ‘Zero on Both Sides’, $\{(1, 1), (8, 4), (4, 8)\}$ for ‘Nonzero on Both Sides’, and $\{(1, 5), (5, 2), (4, 8)\}$ for ‘Spike on Left Side’. When evaluating the mean integrated squared error performance of estimators of the density of $Y(1)$, $Q(0)$ is set equal to $Q(1)$.



(a) Parameter choices leading to improved power against smoother alternatives: $c = 2.5$ and $s = 2$.



(b) Parameter choices leading to improved power against rougher alternatives: $c = 10$ and $s = 0.5$.

Figure S3: Same as Figure 4, but at different choices of the tuning parameters indexing the tests. The tests based on Example 1a set the regularization parameter so that $\beta_k = 1/[1 + (k/c)^2]$ for a constant c and the MMD test selects a bandwidth equal to some constant s times median $\{Y_1, \dots, Y_n\}$. The same ncausal panel from Figure 4 is shown as a benchmark in both subfigures.