

Rejoinder: On nearly assumption-free tests of nominal confidence interval coverage for causal parameters estimated by machine learning

Lin Liu, Rajarshi Mukherjee and James M. Robins

We thank the editors for this opportunity and the discussants Kennedy, Balakrishnan and Wasserman (2020) (abbreviated as KBW in the sequel) for their insightful commentaries on our paper (Liu, Mukherjee and Robins, 2020) (abbreviated as LMR in the sequel).

1. A BRIEF INTRODUCTION TO HIGHER ORDER INFLUENCE FUNCTIONS

We would like to start our rejoinder by responding to the philosophical comments in Section 6 of KBW’s discussion before getting into the other more technical comments. In Section 6, KBW divide statistical procedures into structure-driven and methods-driven but also acknowledge that the boundary between these two categories is blurry. For example, even for the poster child of the methods-driven tools – deep neural networks – one common research direction is to prove some form of optimality or robustness under some assumptions, often quantified by smoothness, sparsity or other related complexity measures such as metric entropy (Schmidt-Hieber, 2020, Hayakawa and Suzuki, 2020, Barron and Klusowski, 2018).

The discussants then state that higher order influence function (HOIF) based methods are ‘structure-driven’ because ‘they typically rely on carefully constructed series estimates’ and achieve ‘better performance over appropriate Hölder spaces potentially at the expense of being more structure driven.’ This statement misunderstands the motivation and goals of HOIF estimation. Our goal has always been to make HOIF fully methods-driven. However,

Lin Liu is Assistant Professor, Institute of Natural Sciences, School of Mathematical Sciences and SJTU-Yale Joint Center for Biostatistics, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: linliu@sjtu.edu.cn). Rajarshi Mukherjee is Assistant Professor, Department of Biostatistics, Harvard T. H. Chan School of Public Health, Boston, Massachusetts 02115, USA (e-mail: ram521@mail.harvard.edu). James M. Robins is Professor, Department of Epidemiology and Biostatistics, Harvard T. H. Chan School of Public Health, Boston, Massachusetts 02115, USA (e-mail: robins@hsph.harvard.edu).

before we reach this goal, difficult open problems remain to be solved. Until then, we have had to make restrictive assumptions to obtain sharp mathematical results – these assumptions can make our methodology appear at least partly ‘structure-driven’.

The theory of HOIF is (simplifying somewhat) a theory based only on higher order scores of finite dimensional submodels. As a consequence, the theory by itself cannot quantify the rates of convergence of a HOIF estimator and thus the bias of a HOIF estimator without additional complexity reducing model assumptions, a central point we stressed throughout LMR. To be more concrete, for now let us restrict the attention to smooth nonlinear functionals $\psi(\theta)$ of a distribution P_θ lying in an infinite dimensional model $\mathcal{M} = \{P_\theta; \theta \in \Theta\}$ with a first order influence function $\mathbb{IF}_{1,\psi}(\theta)$ but (as is generally the case in infinite dimensional models) without m th order influence functions for $m > 1$. Therefore, HOIF theory often considers finite $k = k(n)$ -dimensional sieves $\mathcal{M}_{\text{sub},k} = \{P_\theta; \theta \in \Theta_{\text{sub},k} \subset \Theta\}$ containing an initial training sample estimator $\hat{\theta}$, an associated projection map $\theta \mapsto \tilde{\theta}_k$ from Θ onto $\Theta_{\text{sub},k}$ that is the identity for $\theta \in \Theta_{\text{sub},k}$. The projection map defines a truncated parameter $\tilde{\psi}_k(\theta), \theta \in \Theta$ by $\tilde{\psi}_k(\theta) = \psi(\tilde{\theta}_k(\theta)), \theta \in \Theta$, which will typically have HOIFs of all orders because $\Theta_{\text{sub},k}$ is finite dimensional. The theory of HOIF applied to the parameter $\tilde{\psi}_k(\theta)$ guarantees that $\{\tilde{\psi}_k(\hat{\theta}) + \mathbb{E}_\theta[\mathbb{IF}_{m,\tilde{\psi}_k}(\hat{\theta})]\} - \tilde{\psi}_k(\theta) = O(\|\hat{\theta} - \theta\|^{m+1})$ or, equivalently,

$$\mathbb{E}_\theta[\hat{\psi}_{m,k} - \tilde{\psi}_k(\theta)] \equiv \mathbb{E}B_{\theta,m,k}(\hat{\psi}_1) = O(\|\hat{\theta} - \theta\|^{m+1})$$

where $\hat{\psi}_{m,k} = \psi(\hat{\theta}) + \mathbb{IF}_{m,\tilde{\psi}_k}(\hat{\theta})$.

Here $\hat{\psi}_1 = \psi(\hat{\theta}) + \mathbb{IF}_{1,\tilde{\psi}_k}(\hat{\theta})$ is a doubly robust machine learning (DRML) estimator based on the first order influence function and $\mathbb{IF}_{m,\tilde{\psi}_k}(\hat{\theta}) = \mathbb{IF}_{1,\tilde{\psi}_k}(\hat{\theta}) - \sum_{j=2}^m \mathbb{IF}_{jj,\tilde{\psi}_k}(\hat{\theta})$ where, under $P_{\hat{\theta}}$, $\mathbb{IF}_{jj,\tilde{\psi}_k}(\hat{\theta}) \equiv \widehat{\mathbb{IF}}_{jj,k}$ is a j th order U -statistic.¹ Unless stated otherwise all expectations are conditional on the training sample.

¹Here we are using the same sign convention as in LMR, which reverses the sign conventions of Robins et al. (2008).

Thus $\mathbb{E}_\theta[\hat{\psi}_{m,k} - \psi(\theta)] = \text{EB}_{\theta,m,k}(\hat{\psi}_1) + \text{TB}_{\theta,k}(\hat{\psi}_1)$ with $\text{TB}_{\theta,k}(\hat{\psi}_1) \equiv \tilde{\psi}_k(\theta) - \psi(\theta)$. Furthermore, it is often the case that $\text{var}_\theta(\hat{\psi}_{m,k}) = O(k^{m-1}/n^m \vee 1/n)$. The above is pretty much the cornerstone of the theory of HOIF estimators. This theory involves no structural assumptions on components of θ , such as smooth or sparse nuisance functions. As a consequence the theory is agnostic as to the rate at which $\|\hat{\theta} - \theta\|$ or $\text{TB}_{\theta,k}(\hat{\psi}_1)$ converges to zero.

The above theory was introduced in Sections 2–3 of [Robins et al. \(2008\)](#) before either Hölder smoothness assumptions or best approximating bases were introduced. However, we then went on to study models defined in terms of Hölder smoothness to determine whether our abstract theory (just described) could be used to construct rate minimax estimators (it could) for a particular class of functionals under this well-known infinite dimensional model. Under the Hölder model, we could determine the rate at which $\text{EB}_{\theta,m,k}(\hat{\psi}_1)$ and $\text{TB}_{\theta,k}(\hat{\psi}_1)$ converged to 0 for different choices of the sequences $m = m(n)$, $k = k(n)$, and parametric submodels $\mathcal{M}_{\text{sub},k}$. We could thus optimize $m(n)$, $k(n)$, and $\mathcal{M}_{\text{sub},k(n)}$ and often obtained minimax rates under the Hölder model, when we did so.

Indeed, our theoretical work on HOIF since [Robins et al. \(2008\)](#) and [Robins et al. \(2017\)](#) can be understood as having been solely directed toward the elimination of remaining structure-driven assumptions. As an example, consider the parameter $\psi(\theta) = \mathbb{E}_\theta[\text{var}_\theta[A|X]]$ with A Bernoulli and X high-dimensional with a continuous distribution. Then $\theta = (p, g)$ where $p(x) = \mathbb{E}_\theta[A|X = x]$ and $g(x)$ is the density of X . A k -dimensional submodel $p(x; \theta_k) \equiv p_{\theta_k}(x)$ for $p(x)$ was chosen to be $\{p_{\theta_k}; p_{\theta_k}(x) = \hat{p}(x) + \theta_k^\top \bar{z}_k(x)\}$ where $\bar{z}_k(x)$ is the vector of the first k elements of a sequence of user-selected set of basis functions $\{\bar{z}_j(x), j = 1, \dots\}$. The corresponding projection map is $\tilde{\theta}_k = -\Omega_k^{-1} \mathbb{E}_\theta[A(\hat{p}(X) - p(X))]$ with $\Omega_k := \mathbb{E}_\theta[\bar{z}_k(X)\bar{z}_k(X)^\top] \equiv \mathbb{E}_g[\bar{z}_k(X)\bar{z}_k(X)^\top]$. The HOIFs $\mathbb{IF}_{m,\tilde{\psi}_k}(\theta)$ depended on g through Ω_k which we estimated by $\mathbb{E}_{\hat{g}}[\bar{z}_k(X)\bar{z}_k(X)^\top]$ with \hat{g} an estimator of the density g . In the above papers, we used complexity reducing models (e.g., Hölder models) on both p and g to evaluate the rate of convergence of $\text{EB}_{\theta,m,k}(\hat{\psi}_1) = O(\|\hat{p} - p\|^2 \|\hat{g} - g\|^{m-1})$ to zero. In the case in which $k = o(n)$, [Mukherjee, Newey and Robins, 2017](#) introduced “empirical” HOIF estimators that eliminated the need to assume a complexity-reducing model on g . Instead they proposed estimating $\{\mathbb{E}_g[\bar{z}_k(X)\bar{z}_k(X)^\top]\}^{-1}$ by the inverse sample Gram matrix $\{\Omega_k^{tr}\}^{-1} \equiv \{\mathbb{P}_{n_{tr}}[\bar{z}_k(X)\bar{z}_k(X)^\top]\}^{-1}$ in the training sample for $k < n$. Indeed, the goal of LMR was to determine the inferential questions concerning smooth nonlinear functionals that remain answerable when one refuses to impose any complexity reducing structural assumptions—a goal that seems to us to be extremely “methods-driven.”

However, several difficult open problems remain to be solved before HOIF inference becomes fully methods-driven; i.e. becomes a robust, off-the-shelf, widely applicable methodology for inference on nonlinear functionals in non- and semiparametric models. We have previously discussed these remaining problems both in LMR and earlier papers. In this rejoinder, we discuss some of them in greater depth to respond to discussants’ concerns and suggestions.

2. TOWARDS “METHODS-DRIVEN” HOIFs

The main bottleneck in achieving fully “methods-driven” HOIF inference is the dependence of the power of our falsification tests and the efficiency of our estimators on the choice of the basis functions $\bar{z}_k(x)$. In Section 2.1, we propose a relaxation of one of the assumptions in LMR that dispenses with our reliance on ‘carefully constructed’ choices (such as compactly supported wavelets or B-splines) for the basis functions $\bar{z}_k(x)$, at the cost of perhaps a small loss in power.

In their final section, KBW consider one of the most interesting open problem in the theory of HOIF: how to adaptively select the m basis functions $\tilde{f}_m(x) = (f_1(x), \dots, f_m(x))$ of the d -dimensional vector x to (approximately) minimize the truncation bias for the expected conditional variance.² We have been investigating this same problem for several years but we have yet to come up with a wholly satisfying approach. KBW suggest a new approach based on aggregation. In Section 2.2, we show by a toy example that this approach seems promising and is worth further in-depth investigation. However, we also raise a difficult problem that needs to be solved before this promise can be fulfilled.

2.1 Dispensing with the Need for Carefully Constructed Basis Functions

LMRs assumed Condition W in the statement and proof of Theorems 3.2 and 4.2. Condition W imposes severe restrictions on the basis function \bar{z}_k that can be chosen. Here we show these restrictions can be avoided by replacing Condition W with Condition SW below. In fact, we mentioned Condition SW in Remark 2.5 of LMR but failed to provide sufficient emphasis and context. With the exception of the online supplement, following the recommendation of a referee, LMR focus on the semisupervised

²Following KBW, we have used m rather than k to indicate the dimension of the vector of basis functions in their statistic $\widehat{\mathbb{IF}}_{22,\text{KBW}}$ defined in Section 2.2. Note that by Theorem 3.2 of LMR, we require $m < n$ to have power to reject the null hypothesis under the alternative that $\frac{\text{Bias}_{\theta,k}(\hat{\psi}_1)}{\text{s.e.}_\theta(\hat{\psi}_1)} = \delta + c$ for any given $c > 0$. Recall that our ability to detect with probability going to 1 any alternative of order $c/n^{1/2}$ for any fixed $c > 0$ is one of the perhaps surprising consequences of our tests based on HOIF (due to the fact that with $k < n$, the variance of $\widehat{\mathbb{IF}}_{22,k}$ under that alternative is k/n^2 .)

setting in which $\mathbb{E}_g[\bar{z}_k(X)\bar{z}_k(X)^\top]$ is known. In that case, we only require the following weakened form of Condition W in LMR for the level and power properties stated in Theorems 3.2 and 4.2 of our test $\hat{\chi}_k(z_{\alpha/2}, \delta)$ for the surrogate null hypothesis $H_{0,k}(\delta) : \frac{|\text{Bias}_{\theta,k}(\hat{\psi}_1)|}{\text{s.e.}_{\theta}[\hat{\psi}_1]} \leq \delta$ to hold, where $\text{Bias}_{\theta,k}(\hat{\psi}_1) = \mathbb{E}_{\theta}[\hat{\psi}_1 - \tilde{\psi}_k(\theta)]$.

CONDITION SW.

1. All the eigenvalues of Ω_k are bounded away from 0 and ∞ ;
2. The true nuisance functions $b(X)$ and $p(X)$, and the estimated nuisance functions $\hat{b}(X)$ and $\hat{p}(X)$, are all bounded with \mathbb{P}_{θ} -probability 1;
3. $\|\bar{z}_k(x)^\top \bar{z}_k(x)\|_{\infty} \leq Bk$ for some constant $B > 0$.

Condition SW weakens Condition W in LMR by dropping the requirement that $\|\Pi[\hat{b} - b|\bar{z}_k]\|_{\infty} \leq C$ and $\|\Pi[\hat{p} - p|\bar{z}_k]\|_{\infty} \leq C$ for some constant $C > 0$ not depending on n . This extra condition holds for wavelets, B-spline and local polynomial partition series (Belloni et al., 2015). However, there are many additional choices of \bar{z}_k that satisfy Condition SW without satisfying Condition W, including Fourier series and monomial transformations of the covariates X when X is compactly supported or monomial transformations of some bounded transformation of X when X is unbounded. Allowing \bar{z}_k to include monomial transformations of the covariates makes our approach more flexible and “methods-driven.”

Turn now to the case considered in the supplement of LMR and Liu, Mukherjee and Robins (2020a) in which the expected Gram matrix $\Omega_k = \mathbb{E}_g[\bar{z}_k(X)\bar{z}_k(X)^\top]$ is unknown and therefore tests of $H_{0,k}(\delta)$ must now be based on empirical HOIFs that substitute $\hat{\Omega}_k = \mathbb{P}_{n_{tr}}[\bar{z}_k(X)\bar{z}_k(X)^\top]$ for Ω_k . In Section S1 (also in Section S3 of the supplement of LMR and Liu, Mukherjee and Robins (2020a), Sections 4 and 5), we show in that case that the additional conditions $\|\Pi[\hat{b} - b|\bar{z}_k]\|_{\infty} \leq C$ and $\|\Pi[\hat{p} - p|\bar{z}_k]\|_{\infty} \leq C$ are needed for tests $\hat{\chi}_{33,k}(\hat{\Omega}_k^{-1}; z_{\alpha/2}, \delta)$ (see equation (S1.1)) that use $\hat{\Omega}_k$ to attain the same asymptotic power as the oracle tests $\hat{\chi}_k(z_{\alpha/2}, \delta)$ that use Ω_k . When these infinity-norm bounds do not hold (e.g., for Fourier series or monomial transformation of compactly-supported covariates), the asymptotic power of the test might be less. However, the level of the test $\hat{\chi}_{33,k}(\hat{\Omega}_k^{-1}; z_{\alpha/2}, \delta)$ under Conditions W and SW are identical under some additional restrictions.³ Hence, at least in the context of bias testing, we shall relax Condition W to Condition SW in the future (e.g., Liu, Mukherjee

and Robins, 2020a), whether or not $\mathbb{E}_g[\bar{z}_k(X)\bar{z}_k(X)^\top]$ is known, so as to remove restriction to “carefully constructed series.” The only cost is a possible small loss in power and that only if the infinity norms of the projections $\Pi[\hat{b} - b|\bar{z}_k]$ and/or $\Pi[\hat{p} - p|\bar{z}_k]$ are not bounded even though those of $\hat{b} - b$ and $\hat{p} - p$ are bounded under Condition SW. In summary, by adopting Condition SW rather than Condition W, our methodology becomes more “methods-driven” and allows \bar{z}_k to include most basis functions relevant for practice.

2.2 KBW’s Aggregation Approach

In the following, to avoid irrelevant issues, we consider the case in which the density g of X is known, $\psi(\theta) = \mathbb{E}_{\theta}[\text{var}_{\theta}[A|X]]$, and our inferential goal is to test the actual null hypothesis $H_0(\delta) : \text{Bias}_{\theta}(\hat{\psi}_1) \leq \delta \text{s.e.}_{\theta}(\hat{\psi}_1)$ as in LMR.

KBW propose the following procedure. First divide the data into three randomly selected subsamples: a training sample \mathcal{D}_{tr} , a selection (auxiliary) sample \mathcal{D}_{sel} , and an estimation sample \mathcal{D}_{est} . We are given an estimate $\hat{p}(x)$ of $\mathbb{E}_{\theta}[A|X = x]$ obtained from \mathcal{D}_{tr} . We then use data \mathcal{D}_{sel} to regress the residuals $A - \hat{p}(X)$ using m different methods to obtain $\hat{f}(x) = \{\hat{f}_{\ell}(x), \ell = 1, \dots, m\}$ predictors of the true residual function $p(x) - \hat{p}(x)$. Finally in sample \mathcal{D}_{est} , we compute⁴

$$(2.1) \quad \begin{aligned} & \widehat{\mathbb{I}\mathbb{F}}_{22,KBW}(\hat{f}_m) \\ &= \frac{1}{n(n-1)} \sum_{i_1 \neq i_2 \in \mathcal{D}_{est}} (A_{i_1} - \hat{p}(X_{i_1})) \hat{f}_m(X_{i_1})^\top \\ & \quad \times \Omega_{\hat{f}_m}^{-1} \hat{f}_m(X_{i_2}) (A_{i_2} - \hat{p}(X_{i_2})). \end{aligned}$$

where $\Omega_{\hat{f}_m} \equiv \mathbb{E}_g[\hat{f}_m(X)\hat{f}_m(X)^\top | \mathcal{D}_{sel}] = \int \hat{f}_m(x) \times \hat{f}_m(x)^\top g(x) dx$. KBW test the hypothesis $H_0(\delta)$, using their test statistic

$$\begin{aligned} & \widehat{\chi}_{KBW,m}(z_{\alpha/2}, \delta) \\ &= \mathbb{1} \left\{ \frac{\widehat{\mathbb{I}\mathbb{F}}_{22,KBW}(\hat{f}_m)}{\widehat{\text{s.e.}}(\hat{\psi}_1)} - z_{\alpha/2} \frac{\widehat{\text{s.e.}}(\widehat{\mathbb{I}\mathbb{F}}_{22,KBW}(\hat{f}_m))}{\widehat{\text{s.e.}}(\hat{\psi}_1)} > \delta \right\} \end{aligned}$$

which is precisely the test statistic $\hat{\chi}_k(z_{\alpha/2}, \delta)$ in Section 3 of LMR except that $\widehat{\mathbb{I}\mathbb{F}}_{22,k}$ is replaced by $\widehat{\mathbb{I}\mathbb{F}}_{22,KBW}(\hat{f}_m)$.

Roughly speaking, KBW’s idea is that if one among the m ($m < n$) methods both captures the true structure class (e.g., a smoothness class versus a sparsity class) of the true residual function $p(x) - \hat{p}(x)$ and appropriately chosen tuning parameters, then the power of the KBW test $\widehat{\chi}_{KBW,m}(z_{\alpha/2}, \delta)$ will be equal to or greater than the power

³For the additional restrictions, see Proposition S1.1 and Remark S3.5 of the supplement of LMR. In Liu, Mukherjee and Robins (2020a), we show that it is possible to remove these *additional restrictions* by extending U -statistics of order three to diverging order with increased computational cost.

⁴As in LMR, we always condition on \mathcal{D}_{tr} , which is therefore suppressed in the notation.

of the LMR test $\widehat{\chi}_k(z_{\alpha/2}, \delta)$ that would be used by an oracle who knew the true structure class and the optimal basis functions $\bar{z}_k(x)$ that maximize the power of $\widehat{\chi}_k(z_{\alpha/2}, \delta)$.

Below we will show how the KBW test performs in a toy example under Hölder assumptions.

EXAMPLE. Let X be one-dimensional, $p(x)$ be Hölder with exponent s with $s < 1/2$ and $\hat{p}(x)$ be a series estimator of p using \bar{z}_{k^*} , the first $k^* \asymp n^{1/(1+2s)}$ compactly supported Cohen–Daubechies–Vial (CDV) wavelets (of sufficient order).

This k^* is minimax optimal for estimating p under mean squared error loss. With this minimax optimal \hat{p} , $\text{Bias}_{\theta}(\hat{\psi}_1) = \mathbb{E}_{\theta}[(p(X) - \hat{p}(X))^2] \asymp n^{-2s/(1+2s)}$. Since $s < 1/2$, $\text{Bias}_{\theta}(\hat{\psi}_1) > n^{-1/2}$ and thus $H_0(\delta)$ is false. For any k such that $n > k \gg k^*$, $\text{Bias}_{\theta,k}(\hat{\psi}_1) \equiv \mathbb{E}_{\theta}[\widehat{\mathbb{I}}\mathbb{F}_{22,k}] \asymp \text{Bias}_{\theta}(\hat{\psi}_1)$ (or equivalently $\frac{\text{TB}_{\theta,k}(\hat{\psi}_1)}{\text{Bias}_{\theta}(\hat{\psi}_1)} \rightarrow 0$) and by Theorem 3.2 of LMR, the LMR test $\widehat{\chi}_k(z_{\alpha/2}, \delta)$ rejects $H_0(\delta)$ with probability approaching 1 with increasing sample size.

In the same setup, we consider the KBW test $\widehat{\chi}_{\text{KBW},m}(z_{\alpha/2}, \delta)$ with $m = 1$ and \hat{f} the estimated regression function of $A - \hat{p}(X)$ on $\bar{z}_k(X)$, with $k^* \ll k \ll (\frac{k^{*2}}{\sqrt{n}} \wedge n)$. Below we argue that the KBW test $\widehat{\chi}_{\text{KBW},m=1}(z_{\alpha/2}, \delta)$ with such \bar{z}_k rejects $H_0(\delta)$ with probability approaching 1. The reason for $k \ll \frac{k^{*2}}{\sqrt{n}}$ will be explained in the end of this section. Thus in summary, for every $s \in (0, 1/2)$, we can construct a test $\widehat{\chi}_k(z_{\alpha/2}, \delta)$ as in LMR and a test $\widehat{\chi}_{\text{KBW},m=1}(z_{\alpha/2}, \delta)$ as in KBW to reject $H_0(\delta)$ with probability approaching 1.

We now complicate our example by supposing that we only know $s \in (0, 1/2)$ but do not know its true value. Therefore, \hat{p} will be an optimal adaptive estimator of p estimated from the training sample \mathcal{D}_{tr} . Our goal now is to construct an adaptive test that rejects with probability going to 1 whatever the true value s is. The LMR test $\widehat{\chi}_k(z_{\alpha/2}, \delta)$ with $k = n/c$ with any $c > 1$ is adaptive in this sense. However, no KBW test $\widehat{\chi}_{\text{KBW},m=1}(z_{\alpha/2}, \delta)$ can be adaptive. To see this, if $s = 1/2 - \epsilon$ for some very small $\epsilon > 0$, then the value of k used for the test $\widehat{\chi}_{\text{KBW},m=1}(z_{\alpha/2}, \delta)$ must be less than $\frac{k^{*2}}{\sqrt{n}} = n^{\frac{1}{2} + \epsilon}$ with k^* evaluated at $n^{\frac{1}{1+2(1/2-\epsilon)}}$. On the other hand, if $s = \epsilon$, k must be larger than $k^* = n^{\frac{1}{1+2\epsilon}}$. Hence, there exists no single test $\widehat{\chi}_{\text{KBW},m=1}(z_{\alpha/2}, \delta)$ that can adapt. It remains an open question whether there exists an adaptive test $\widehat{\chi}_{\text{KBW},m}(z_{\alpha/2}, \delta)$ for any choice of m , even in this simple example.

We now justify the above claims. KBW first estimate the regression of $A - \hat{p}(X)$ on $\bar{z}_k(X)$ from the selection sample \mathcal{D}_{sel} as follows:

$$\hat{f}(x) = \bar{z}_k(x)^{\top} \hat{\beta},$$

where $\hat{\beta} = \Omega_k^{-1} \frac{1}{n} \sum_{i \in \mathcal{D}_{\text{sel}}} \bar{z}_k(X_i)(A_i - \hat{p}(X_i))$. Then KBW compute $\widehat{\mathbb{I}}\mathbb{F}_{22,\text{KBW}}(\hat{f})$ as in equation (2.1).

Recall that by choosing $k \gg k^*$, $\text{Bias}_{\theta,k}(\hat{\psi}_1) \asymp \text{Bias}_{\theta}(\hat{\psi}_1)$. Following the proof of Theorem 3.2 of LMR, for $\widehat{\chi}_{\text{KBW},1}(z_{\alpha/2}, \delta)$ to reject the null hypothesis $H_0(\delta)$ with probability approaching 1, we need:

- (i) the mean of $\widehat{\mathbb{I}}\mathbb{F}_{22,\text{KBW}}(\hat{f})$ exceeds $n^{-1/2}$: $\mathbb{E}_{\theta}[\widehat{\mathbb{I}}\mathbb{F}_{22,\text{KBW}}(\hat{f})] \asymp \frac{\text{Bias}_{\theta,k}(\hat{\psi}_1)^2}{k/n} \gg \text{s.e.}_{\theta}(\hat{\psi}_1) \asymp \frac{1}{\sqrt{n}}$; and
- (ii) the mean of $\widehat{\mathbb{I}}\mathbb{F}_{22,\text{KBW}}(\hat{f})$ dominates its standard error: $\mathbb{E}_{\theta}[\widehat{\mathbb{I}}\mathbb{F}_{22,\text{KBW}}(\hat{f})] \gg \text{s.e.}_{\theta}[\widehat{\mathbb{I}}\mathbb{F}_{22,\text{KBW}}(\hat{f})]$.

In Section S2, we show that $\widehat{\mathbb{I}}\mathbb{F}_{22,\text{KBW}}(\hat{f})$ has mean of order $\frac{\text{Bias}_{\theta,k}(\hat{\psi}_1)^2}{k/n}$, which dominates its standard error of order $\frac{\text{Bias}_{\theta,k}(\hat{\psi}_1)/\sqrt{n}}{k/n}$ when $H_0(\delta)$ is false and $\text{Bias}_{\theta,k}(\hat{\psi}_1) \asymp \text{Bias}_{\theta}(\hat{\psi}_1)$. Hence, (ii) should hold. In terms of (i), if $H_0(\delta)$ is false, that is, $\text{Bias}_{\theta}(\hat{\psi}_1) \asymp \frac{k^*}{n} \gg \frac{1}{\sqrt{n}}$ or equivalently $k^* \gg \sqrt{n}$ (or equivalently the smoothness index $s < 1/2$), there always exists $k \gg k^*$ also satisfying (i) as (1) $k^* \ll \frac{k^{*2}}{\sqrt{n}}$ when $k^* \gg \sqrt{n}$ and (2) (i) is equivalent to

$$\frac{\text{Bias}_{\theta,k}(\hat{\psi}_1)^2}{k/n} \asymp \frac{\text{Bias}_{\theta}(\hat{\psi}_1)^2}{k/n} \asymp \frac{k^{*2}}{kn} \gg \frac{1}{\sqrt{n}} \Leftrightarrow k \ll \frac{k^{*2}}{\sqrt{n}}.$$

3. CLASSES OF FUNCTIONALS AND THE BIAS TEST

3.1 Monotone Bias Class and Alternative Sample Splitting Schemes

The analysis of the expected density example described in Section 2 of KBW illustrates our approach in a simple setting and is essentially isomorphic to the analysis of the expected conditional variance $\mathbb{E}_{\theta}[\text{var}_{\theta}[A|X]]$ when, as in the semisupervised setting considered in the main text of LMR, the marginal distribution of X is known. Both of these functionals are members of our monotone bias class, a class so named because of our claim that the bias of the our second order U -statistic estimator was nonincreasing in the number of basis functions k . KBW show that our claim does not hold if, following Newey and Robins (2018), one uses two different estimates of the density of X (or of $\mathbb{E}_{\theta}[A|X]$ in the $\mathbb{E}_{\theta}[\text{var}_{\theta}[A|X]]$ example) coming from two independent subsamples of the training sample. This is a fact we were well aware of (see Section S.1.1 of the supplement of LMR).

The multiple training sample splitting nuisance function estimators of Newey and Robins (2018) can sometimes achieve faster convergence rates than the (single training sample) DRML estimators considered in the main text of LMR. In fact, they can even achieve $n^{-1/2}$ rates for estimation of $\mathbb{E}_{\theta}[\text{var}_{\theta}[A|X]]$ and $\mathbb{E}_{\theta}[\text{cov}_{\theta}[A, Y|X]]$ under minimal Hölder smoothness assumptions (Robins et al., 2009) needed for \sqrt{n} -estimation. But this result requires, not only that the Hölder class assumptions are true,

but also that one uses particular *undersmoothed* nuisance function estimators (such as local polynomials or regression splines) rather than flexible black-box machine learning (such as random forests or deep learning) estimators, that are the motivation for and the subject of our paper.

Indeed, it was our intention to define DRML estimators to be estimators in which all nuisance functions were estimated with ML algorithms from the same training sample, as this is the current ‘state of the art’; and, as we emphasized in the Introduction of our paper, our approach is one of being “in dialogue with current practice and practitioners.” In fact, our technical Lemma 2.3 of LMR, which is the lemma in which we differentiate the bias properties of $\mathbb{E}_\theta[\text{var}_\theta[A|X]]$ from those of $\mathbb{E}_\theta[\text{cov}_\theta[A, Y|X]]$, is correct as stated, as the lemma explicitly applies to the case in which a single estimator of \hat{p} of $\mathbb{E}_\theta[A|X]$ is used in the DRML estimator $\hat{\psi}_1$ of $\mathbb{E}_\theta[\text{var}_\theta[A|X]]$ (as is evident from the paragraph just prior to the Lemma).

There remain important open problems that we, with other colleagues, and our discussants are considering in ongoing work on the properties of multiple training sample splitting nuisance function estimators: (1) is it possible to develop a general theory of the benefit of multiple training sample splitting, even when black-box machine learning estimators of the nuisance functions are employed and (2) can undersmoothing be automated to directly balance bias with variance for the estimators of the functional of interest? The model selection strategy in Cui and Tchetgen Tchetgen (2019) may be a viable option.

3.2 Generalization to Functionals with the Mixed Bias Property

In Liu, Mukherjee and Robins (2020a), we also show that the methods of LMR and the theory of HOIFs (Robins et al., 2008, 2016) can be extended to the entire class of parameters/functionals with the so-called *mixed bias property* (henceforth called MB functionals), studied by Rotnitzky, Smucler and Robins (2019). This class is a strict superset of the union of two overlapping classes of functionals introduced in Robins et al. (2008) and Chernozhukov, Newey and Singh (2018), respectively.

DEFINITION 3.1 (Definition 1 of Rotnitzky, Smucler and Robins, 2019). A parameter/functional $\psi(\theta)$ is an MB functional if, for each $\theta \in \Theta$, there exist $b : x \mapsto b(x) \in \mathcal{B}$ and $p : x \mapsto p(x) \in \mathcal{P}$ such that (i) $\theta = (b, p, \theta_{\setminus(b,p)})$ and $\Theta = \mathcal{B} \times \mathcal{P} \times \Theta_{\setminus(\mathcal{B},\mathcal{P})}$ and (ii) for any θ, θ'

$$(3.1) \quad \begin{aligned} & \psi(\theta') - \psi(\theta) + \mathbb{E}_\theta[\text{IF}_1(\theta')] \\ & = \mathbb{E}_\theta[S_{bp}(b(X) - b'(X))(p(X) - p'(X))], \end{aligned}$$

where $S_{bp} \equiv s_{bp}(O)$ and $o \mapsto s_{bp}(o)$ is a known function that does not depend on θ or θ' satisfying either $\mathbf{P}_\theta(S_{bp} \geq 0) = 1$ or $\mathbf{P}_\theta(S_{bp} \leq 0) = 1$ and $\text{IF}_1(\theta)$ is the (first order) influence function of the parameter $\psi(\theta)$.

4. ON THE COVARIATE STRUCTURE

KBW raise the interesting question of “how and whether randomness of the covariates and/or smoothness of the covariate density should be relied on in practice.” They suggest that, perhaps, the covariates X should be conditioned on (and thus be regarded as fixed rather than random) in any inferential procedure, whenever study subjects are not randomly sampled from some well-defined population. If this suggestion were followed the benefits of HOIFs may be greatly diminished, but, if so, HOIFs would not be alone. In many areas, it is essential that the covariates X and the outcome Y have a joint distribution (X, Y) – conformal inference, prediction risk minimization, the bias-corrected lasso, and covariate shift methods being four current examples, at least two of which our discussants have written about. In fact, if the subjects were not randomly sampled, one could equally ask why should we consider Y conditioned on $X = x$ as random, as we have no reason to believe any measure of their association is invariant across studies or populations, especially if noncausal.

So the question is how to inject randomness into an observational study, a question that arises when an investigator wishes to generalize her findings from the observed study subjects to some larger population. For example, an investigator who considers recommending a public health intervention, would hope to have studied subjects that are in some sense representative of the population of potential recipients. The simplest random model allowing generalization is to consider the study subjects as a random sample from some very large (effectively infinite) hypothetical (i.e., fictitious) superpopulation of potential recipients with the superpopulation empirical distribution serving as the target of inference. This is effectively reverse engineering, in the sense that even if one has studied a ‘convenience sample’, one can still hypothesize a superpopulation that is similar, allows the use of ordinary i.i.d. statistical methods, to obtain valid confidence intervals for functionals of the superpopulation empirical distribution (see Robins, 1988). This approach may seem distasteful (or even vacuous) to a purist and not as elegant as de Finetti’s subjectivist approach, but we believe it underlies what frequentist analyses that epidemiologists and statisticians are doing daily, perhaps without explicit recognition. Of course, generalization from a convenience sample to an actual, nonhypothetical, population is possible only based on further substantive knowledge. Typical examples of the sort of convenience samples we are thinking of include (i) all members of a HMO admitted for an acute myocardial infarction between Jan 1, 2015 and Dec 31, 2017 or (ii) all workers employed at a particular asbestos mine at any point in the interval 1959 to 1965.

Putting such philosophical matters aside, let us for now assume X is random. In this case, KBW raised the question whether the performance of HOIF tests and estimators is robust to misspecification of a model for the density g of X say in the estimation of $\mathbb{E}_\theta[\text{var}_\theta[A|X]]$. The answer is HOIFs are not robust as can be seen from the fact that the estimation bias $\mathbb{E}_\theta[\hat{\psi}_{m,k} - \tilde{\psi}_k(\theta)] \equiv \text{EB}_{\theta,m,k}(\hat{\psi}_1) = O(\|\hat{p} - p\|^2 \|\hat{g} - g\|^{m-1})$. Thus, if \hat{g} is inconsistent with $\|\hat{g} - g\| > 1$, then $\hat{\psi}_{m,k}$ may have bias even greater than the bias $O(\|\hat{p} - p\|^2)$ of the DRML estimator $\hat{\psi}_1$. If, following KBW's suggestion, we react to this nonrobustness by choosing to make no assumptions whatsoever regarding the density of X , then we need to restrict ourselves to the use of empirical higher order influence functions, as they do not require that we estimate g (Mukherjee, Newey and Robins, 2017). The following example helps understand their statistical properties.

Following KBW, suppose the functional of interest is $\psi(\theta) = \mathbb{E}_\theta[b_1(X)] = \mathbb{E}_\theta[\text{Ap}(X)Y]$, where $\psi(\theta)$ is equal to the counterfactual mean of an outcome Y when a binary treatment A is set to $a = 1$ under ignorability of treatment A conditional on the d -dimensional covariate X . Here $b_a(x) = \mathbb{E}_\theta[Y|A = a, X = x]$ and $p(x) = 1/\mathbb{E}_\theta[A|X = x]$. Suppose we have a correct Hölder smoothness model for $b_1(x)$ and $p(x)$ having unknown smoothness exponents s_b and s_p unrestricted except for the requirement that $s/d \equiv (s_b + s_p)/(2d) < 1/4$, implying $\psi(\theta)$ cannot be estimated at rate $n^{-1/2}$ (Robins et al., 2009). Suppose however we make no complexity reducing assumptions about $f_X(x)$, the density of X . In that case we would need to use an empirical HOIF estimator $\hat{\psi}_{m,k}^{\text{emp}} = \psi(\hat{\theta}) + \mathbb{I}\mathbb{F}_{m,\tilde{\psi}_k}^{\text{emp}}(\hat{\theta})$ with $\tilde{z}_k(x)$ chosen as appropriate d -dimensional compactly supported CDV wavelets for optimal approximation with $k < n_{tr} = n/2$,⁵ so the inverse of empirical Gram matrix $\mathbb{P}_{n_{tr}}[\tilde{z}_k(X)\tilde{z}_k(X)^\top]$ exists with high probability. It follows from a slight modification of Theorem 5 of Mukherjee, Newey and Robins (2017) that with $k = n/(\log(n))^3$ and $m = \sqrt{\log(n)}$, $\hat{\psi}_{m,k}^{\text{emp}}$ has, up to log terms, truncation bias $n^{-2s/d}$ and variance $O(1/n)$ and negligible estimation bias and thus attains (up to log terms), the rate of convergence $n^{-2s/d}$ found by Wang et al. (2008) for $\mathbb{E}_\theta[\text{var}_\theta[Y|X]]$ under a fixed design. We believe that no other estimator of the counterfactual mean is known to achieve this rate of convergence for all s_b and s_p satisfying $(s_b + s_p)/(2d) < 1/4$ without imposing further assumptions on the density of X . Finally, we conjecture the log terms may be eliminated by decreasing the rate at which k/n converges to zero and appropriately adjusting the rate at which $m(n)$ increases with n .

Next, we turn to KBW's question concerning the performance of HOIF when X is random but we perform inference conditional on X . We begin by providing an example in which unconditional and conditional inference

are essentially equivalent. Robins et al. (2008) consider the question of whether, for estimation of a conditional variance, random regressors provide faster rates of convergence than do fixed equal-spaced regressors, and, if so, how? They consider a setting in which n i.i.d. copies of (Y, X) are observed with X a d -dimensional random vector, with density $f(\cdot)$ bounded away from 0 and infinity and absolutely continuous w.r.t. the uniform measure on the unit cube $[0, 1]^d$. The regression function $b(x) = \mathbb{E}_\theta[Y|X = x]$ is assumed to lie in a given Hölder ball with Hölder exponent $s < 1$. The goal is to estimate $\mathbb{E}_\theta[\text{var}_\theta[Y|X]]$ under the homoscedastic semiparametric model $\text{var}_\theta[Y|X] = \sigma^2$. Under this model, the authors construct a simple estimator $\hat{\sigma}^2$ that converges at rate $n^{-\frac{4s/d}{1+4s/d}}$, when $s/d < 1/4$. Shen et al. (2019) recently proved this estimator was minimax optimal by proving a matching lower bound.⁶

Wang et al. (2008) and Cai, Levine and Wang (2009) earlier proved that if $X_i, i = 1, \dots, n$, are nonrandom but equally spaced in $[0, 1]^d$, the minimax rate of convergence for the estimation of σ^2 is $n^{-2s/d}$ (when $s/d < 1/4$) which is slower than $n^{-\frac{4s/d}{1+4s/d}}$. Thus randomness in X allows for improved convergence rates even though no smoothness assumptions are made regarding $f(\cdot)$.

To explain how this happens, we describe the estimator of Robins et al. (2008). The unit cube in \mathbb{R}^d is divided into $k = k(n) = n^\gamma$, $\gamma > 1$ identical sub-cubes each with edge length $k^{-1/d}$. For each sub-cube with two or more observations, we randomly select two subjects i and j without replacement. We estimate σ^2 in each such sub-cube by $(Y_i - Y_j)^2/2$. Our estimator $\hat{\sigma}^2$ of σ^2 is the average of the sub-cube-specific estimates $(Y_i - Y_j)^2/2$ over all the sub-cubes with at least two observations.

A simple probability calculation shows that the number of sub-cubes containing at least two observations is $O_{P_\theta}(n^2/k)$ so $\hat{\sigma}^2$ has variance of order k/n^2 conditional on $\mathbb{X} = (X_1, \dots, X_n)$.

To compute the conditional bias, observe that for a given sub-cube with i and j selected $\mathbb{E}_\theta[(Y_i - Y_j)^2/2|\mathbb{X}] = \mathbb{E}_\theta[(Y_i - Y_j)^2/2|X_i, X_j] = \sigma^2 + \{b(X_i) - b(X_j)\}^2/2$. However, $|b(X_i) - b(X_j)| = O(\|X_i - X_j\|^s)$ as $s < 1$ and $\|X_i - X_j\| = d^{1/2}O(k^{-1/d})$ as X_i and X_j are in the same sub-cube. It follows that the conditional bias is $\mathbb{E}_\theta[\hat{\sigma}^2 - \sigma^2|\mathbb{X}] = O_{P_\theta}(k^{-2s/d})$. We next find the k that equates variance and squared bias. Specifically, we solve $k/n^2 = k^{-4s/d}$ which gives $k = n^{\frac{2}{1+4s/d}}$. Then the rate of convergence at this optimal k is $n^{-\frac{4s/d}{1+4s/d}}$ conditional on \mathbb{X}

⁶Suppose we change the model by assuming $s > 1$ but in no other way. In that case, when $s/d < 1/4$ no estimator that attains the rate $n^{-\frac{4s/d}{1+4s/d}}$ is known; however no lower bound has been proved that would imply such an estimator is impossible.

⁵ n_{tr} denotes the training sample size.

with high probability, since $n^{-\frac{4s/d}{1+4s/d}}$ is $(k/n^2)^{1/2}$ evaluated at $k = n^{\frac{2}{1+4s/d}}$. But $n^{-\frac{4s/d}{1+4s/d}}$ is also the unconditional rate of convergence since the conditional bias and conditional variance are of order $O(k^{-2s/d})$ and $O(k/n^2)$ with overwhelming probability.⁷

Robins et al. (2008) conclude that the random design estimator has better bias control and hence converges faster than the optimal equal-spaced fixed design estimator, because the random design estimator exploits the $O_{P_\theta}(n^2/n^{\frac{2}{1+4s/d}})$ random fluctuations for which the X 's corresponding to two different observations are only a distance of $O(\{n^{\frac{2}{1+4s/d}}\}^{-1/d})$ apart. In summary, our calculations conditional on \mathbb{X} indicate that the difference in rates is wholly attributable to the difference between the empirical distribution of the $X_i, i = 1, \dots, n$ in a typical realization under the random design and the empirical distribution in the fixed design. That is, given the realized \mathbb{X} , it is of no consequence whether or not it was generated by a random process; all that matters is how the number of bins with at least two observations scales with the total number of bins k .

Now suppose that $\text{var}_\theta[Y|X = x] = \sigma^2(x)$ is heteroscedastic. Then to consider inference conditional on \mathbb{X} we can take $\mathbb{P}_{n_{\text{est}}}\{\text{var}_\theta[Y|X]\} = n^{-1} \sum_{i \in \text{est}} \text{var}_\theta[Y|X_i]$ rather than $\psi(\theta) = \mathbb{E}_\theta[\text{var}_\theta[Y|X]]$ as our object of inference, where we have again randomly split the data into an estimation and a training sample and we note if X is random our object is random. In the random X case, $\mathbb{P}_{n_{\text{est}}}\{\text{var}_\theta[Y|X]\}$ is a $n^{1/2}$ -consistent estimator of $\mathbb{E}_\theta[\text{var}_\theta[Y|X]]$. It is natural to estimate $\mathbb{P}_{n_{\text{est}}}\{\text{var}_\theta[Y|X]\}$ with either the empirical HOIF estimator $\hat{\psi}_{m,k}^{\text{emp}} = \hat{\psi}_{m,k}^{\text{emp}}(\hat{\Omega}_k^{\text{tr}})$ defined earlier or with $\hat{\psi}_{m,k}^{\text{emp}}(\hat{\Omega}_k^{\text{est}})$ where $\hat{\Omega}_k^{\text{est}} = n^{-1} \sum_{i \in \text{est}} \bar{z}_k(X_i) \bar{z}_k(X_i)^\top$ replaces $\hat{\Omega}_k^{\text{tr}}$ in the estimator.⁸ Motivated by KBW's questions concerning conditional inference, it is of great interest to us to study the properties, conditional on \mathbb{X} , of $\hat{\psi}_{m,k}^{\text{emp}}(\hat{\Omega}_k^{\text{tr}})$ and $\hat{\psi}_{m,k}^{\text{emp}}(\hat{\Omega}_k^{\text{est}})$ as estimators of $\mathbb{P}_{n_{\text{est}}}\{\text{var}_\theta[Y|X]\}$.

5. ON UNIVERSAL INFERENCE

Rather than relying on asymptotic theory, Wasserman, Ramdas and Balakrishnan, 2020 construct an ingenious universal confidence set by inverting concentration inequalities of the log-likelihood ratios, hence valid for any

⁷To see this, define the event I_c , for any $c > 0$, as

$$I_c := \left\{ \exists \geq c \frac{n^2}{k} \text{ bins with at least two observations of } X\text{'s} \right\}$$

On this event, we have $(\hat{\sigma} - \sigma)^2 \lesssim n^{-4s/(4s+d)}$ with high probability. Furthermore, for $k \ll n^2$, there exists $c > 0$ such that I_c happens with high probability when the density f_X is bounded away from 0 and ∞ .

⁸ $\mathbb{P}_{22,k}(\hat{\Omega}_k^{\text{est}})$ and $\hat{\psi}_{m,k}^{\text{emp}}(\hat{\Omega}_k^{\text{est}})$ were introduced in Section S3 of the supplement of LMR and its properties were studied via simulation.

sample size n for models with likelihood functions. This is another interesting contribution by two of the discussants to the large body of assumption-free/lean inference literature (e.g., Rinaldo, Wasserman and G'Sell (2019)). In this section, we compare HOIF inference with universal inference for smooth nonlinear functionals. The take-home message is the following:

1. Universal inference requires a likelihood function. For undominated nonparametric problems, KBW consider (i) finite $k = k(n)$ -dimensional sieves $\mathcal{M}_{\text{sub},k} = \{\mathbb{P}_\theta; \theta \in \Theta_{\text{sub},k} \subset \Theta\}$ of increasing dimension k , (ii) (under the conditions of Proposition 7 in Wasserman, Ramdas and Balakrishnan (2020)) an associated projection map $\theta \mapsto \tilde{\theta}_{\text{KL},k}$ from Θ onto $\Theta_{\text{sub},k}$ that is the identity if $\theta \in \Theta_{\text{sub},k}$ and otherwise is the minimizer in KL-divergence between θ and $\Theta_{\text{sub},k}$, and (iii) a preliminary estimator $\hat{\theta}_{\text{KL},k}$ for $\tilde{\theta}_{\text{KL},k}$ from a split sample. When a functional $\psi(\theta)$ is the object of inference, KBW propose to construct a universal confidence interval for the KL-projection parameter $\tilde{\psi}_{\text{KL},k}(\theta) = \psi(\tilde{\theta}_{\text{KL},k}(\theta))$. Universal inference and HOIF inference may choose the same sieve. In that case, under certain laws \mathbb{P}_θ , the KL projection $\tilde{\psi}_{\text{KL},k}(\theta)$ may be equal to the truncated parameter $\tilde{\psi}_k(\theta)$, as we will show below. Without further complexity reducing assumption to quantify the distance between the sieve and the true law \mathbb{P}_θ , inference for $\tilde{\psi}_{\text{KL},k}(\theta)$ in universal inference or $\tilde{\psi}_k(\theta)$ in HOIF inference is the best one could hope for.

2. Universal confidence intervals are guaranteed to cover $\tilde{\psi}_{\text{KL},k}(\theta)$ at the nominal rate for any sample size n under the conditions of Proposition 7 in Wasserman, Ramdas and Balakrishnan (2020). However, we will show that the length of the confidence interval is of order $\|\hat{\theta}_{\text{KL},k} - \tilde{\theta}_{\text{KL},k}\|$,⁹ which is generically $\gg n^{-1/2}$ when $n > k \gg n^{1/2}$. In contrast, HOIF Wald confidence intervals for $\tilde{\psi}_k(\theta)$ centered at $\hat{\psi}_{m,k}^{\text{emp}}$ have length of order $n^{-1/2}$ because estimators $\hat{\psi}_{m,k}^{\text{emp}}$ of $\tilde{\psi}_k(\theta)$ typically have variance of order $1/n$ for $k < n$ and bias less than $n^{-1/2}$ when m is sufficiently large. When $\tilde{\psi}_{\text{KL},k}(\theta) = \tilde{\psi}_k(\theta)$, it follows that HOIF confidence intervals will be narrower than universal confidence intervals; however, nominal coverage of these HOIF Wald confidence intervals for $\tilde{\psi}_k(\theta)$ is guaranteed only in large samples. It is an interesting open problem to construct universal intervals with guaranteed finite sample coverage of optimal length in large samples.¹⁰

We now explain the above statements. To be concrete, we consider the following data generating process and

⁹Here $\|\cdot\|$ is the ℓ_2 norm of a k -dimensional vector.

¹⁰First order influence function based intervals will enjoy the same properties as the higher order intervals when the test of the null hypothesis $H_{0,k}(\delta)$ fails to reject in large samples.

functional of interest to simplify our analysis. We observe n i.i.d. copies of (A, X) , with $X \sim \text{Uniform}([0, 1])$ and $A \sim N(p(X), 1)$ with $p(x) \in L_2([0, 1])$. The goal is to estimate $\psi(\theta) = \mathbb{E}_\theta[p(X)^2]$. Let $\hat{p}(x)$ denote some initial machine learning estimator of the regression function $p(x)$, computed from an independent training sample treated as fixed. Again, we assume that the density of X is known to focus on the important issues.

5.1 Nonparametric Universal Inference

We choose the following sieve for universal inference:

$$\mathcal{M}_{\text{sub},k} = \{N(p_{\theta_k}(X) = \hat{p}(X) + \theta_k^\top \bar{z}_k(X), 1); \theta_k \in \Theta_{\text{sub},k}\}$$

for some $k \equiv k(n)$. The KL-divergence between any member in this sieve $\mathcal{M}_{\text{sub},k}$ and the true law $A \sim N(p(X), 1)$ is

$$\begin{aligned} \text{KL}(p, p_{\theta_k}) &= \mathbb{E}_\theta[(p(X) - p_{\theta_k}(X))^2] \\ &= \mathbb{E}_\theta[(p(X) - \hat{p}(X) - \theta_k^\top \bar{z}_k(X))^2]. \end{aligned}$$

By definition, $\tilde{\theta}_{\text{KL},k}$ minimizes $\text{KL}(p, p_{\theta_k})$ and hence $\tilde{\theta}_{\text{KL},k} = -\Omega_k^{-1} \mathbb{E}_\theta[\bar{z}_k(X)(\hat{p}(X) - p(X))]$. By Proposition 7 in Wasserman, Ramdas and Balakrishnan (2020), a nominal $1 - \alpha$ universal confidence set always covers $\tilde{\theta}_{\text{KL},k}$ with probability at least $1 - \alpha$.

Based on the sieve chosen above, $\tilde{\theta}_{\text{KL},k} = \tilde{\theta}_k$ (see Section 1), and therefore $\tilde{\psi}_{\text{KL},k}(\theta) \equiv \int p_{\tilde{\theta}_{\text{KL},k}}(x)^2 dx \equiv \int p_{\tilde{\theta}_k}(x)^2 dx \equiv \tilde{\psi}_k(\theta)$. This is not surprising because the KL-divergence between two normals is a quadratic form. Such isomorphism breaks down if $A \sim \text{Bernoulli}(p(X))$ even with the same perturbation p_{θ_k} . However, we can easily restore the isomorphism by replacing KL-divergence with χ^2 -divergence. See Section S3.1 for more detail. Thus it will be interesting to generalize Proposition 7 of Wasserman, Ramdas and Balakrishnan (2020) from KL projection to projection based on general f-divergences (Csiszár, 1963, Ali and Silvey, 1966), which include KL- and χ^2 -divergences as special cases.

5.2 On the Length of Universal vs. HOIF Confidence Intervals

In this section, we suppose that $\tilde{\psi}_{\text{KL},k}(\theta)$ in universal inference and $\tilde{\psi}_k(\theta)$ in HOIF inference coincide. It is then natural to compare the length of the confidence intervals for $\tilde{\psi}_{\text{KL},k}(\theta)$ based on these two approaches.

Universal inference first estimates $\tilde{\theta}_{\text{KL},k}$ from half of the sample \mathcal{D}_1 of size $n/2$ by $\hat{\theta}_{\text{KL},k,\mathcal{D}_1}$.¹¹ Switching to another half of the sample \mathcal{D}_2 also of size $n/2$, universal inference

¹¹ $\hat{\theta}_{\text{KL},k,\mathcal{D}_1} = -\hat{\Omega}_{\mathcal{D}_1,k}^{-1} \mathbb{P}_{n,\mathcal{D}_1}[\bar{z}_k(X)(\hat{p}(X) - A)]$ with $\hat{\Omega}_{\mathcal{D}_1,k}^{-1} := [\mathbb{P}_{n,\mathcal{D}_1}[\bar{z}_k(X_i)\bar{z}_k(X_i)^\top]]^{-1}$ where $\mathbb{P}_{n,\mathcal{D}}$ denotes the empirical measure over the sample \mathcal{D} .

first finds a nominal $1 - \alpha$ confidence set $\hat{\Theta}_{\mathcal{D}_2}(\alpha)$ (see equation (S3.3)) for $\tilde{\theta}_{\text{KL},k}$.

To construct a universal confidence interval for the functional $\tilde{\psi}_{\text{KL},k}(\theta)$ based on the confidence set $\hat{\Theta}_{\mathcal{D}_2}(\alpha)$, Wasserman, Ramdas and Balakrishnan (2020) suggest to use the profile universal confidence interval $\hat{\Psi}_{\mathcal{D}_2}^{\text{profile}}(\alpha)$ or the plug-in universal confidence interval $\hat{\Psi}_{\mathcal{D}_2}^{\text{plug-in}}(\alpha)$ (see equations (S3.2) or (S3.1)). In Section S3.2, we show that both intervals have to contain the ‘‘plug-in’’ estimator $\int p_{\hat{\theta}_{\text{KL},k,\mathcal{D}_1}}(x)^2 dx$. Combined with the following lemma¹², a lower bound is obtained on the expected length of both universal confidence intervals:

LEMMA 5.1. *For any confidence interval Ψ containing $\int p_{\hat{\theta}_{\text{KL},k,\mathcal{D}_1}}(x)^2 dx$, if it covers the target parameter $\tilde{\psi}_{\text{KL},k}(\theta)$ with probability at least $1 - \alpha$, then*

$$\begin{aligned} \mathbb{E}_\theta[L(\Psi)|\mathcal{D}_1] &\geq (1 - \alpha) \left| \int p_{\hat{\theta}_{\text{KL},k,\mathcal{D}_1}}(x)^2 - p_{\tilde{\theta}_{\text{KL},k}}(x)^2 dx \right| \\ &\equiv (1 - \alpha) \\ &\quad \times |(\hat{\theta}_{\text{KL},k,\mathcal{D}_1} - \theta_{\text{KL},\mathcal{D}_1})^\top \Omega_k (\hat{\theta}_{\text{KL},k,\mathcal{D}_1} + \theta_{\text{KL},\mathcal{D}_1})|. \end{aligned}$$

Note that the lower bound given in Lemma 5.1 is typically of the same order as $\|\hat{\theta}_{\text{KL},k,\mathcal{D}_1} - \tilde{\theta}_{\text{KL},k}\|$. The length of a plug-in or a profile universal confidence interval is hence of order $\|\hat{\theta}_{\text{KL},k,\mathcal{D}_1} - \tilde{\theta}_{\text{KL},k}\|$, which is typically of order $(k/n)^{1/2} \gg n^{-1/2}$ when $k \gg n^{1/2}$.

In contrast, HOIF estimators unbiasedly estimate $\tilde{\psi}_k(\theta) \equiv \tilde{\psi}_{\text{KL},k}(\theta)$ using a second order U -statistic $\hat{\psi}_{2,k}$ ¹³ with standard error of order $(1/n)^{1/2} \vee (k/n^2)^{1/2}$. Then a large sample HOIF Wald confidence interval $\hat{\psi}_{2,k} \pm z_{\alpha/2} \widehat{\text{se}}(\hat{\psi}_{2,k})$ typically has length of order $n^{-1/2}$ even if $k \gg n^{1/2}$ as long as $k < n$. Even if $k > n$, the length (of order $(k/n^2)^{1/2}$) of an HOIF interval is still shorter than that of a universal confidence interval (of order $(k/n)^{1/2}$).

REMARK 5.2. In Murphy and van der Vaart (2000), the authors (MvdV) showed that, under certain regularity conditions, the confidence interval for $\tilde{\psi}_{\text{KL},k}(\theta)$ based on inverting the profile likelihood ratio test of the hypothesis $\psi = \tilde{\psi}_{\text{KL},k}(\theta)$ attains nominal coverage in large samples and has length of order $n^{-1/2}$ even when $\tilde{\theta}_{\text{KL},k}$ cannot be estimated at rate $n^{-1/2}$. In contrast, as just shown, the universal confidence interval will not shrink at rate $n^{-1/2}$ in this setting; yet, if the universal confidence interval uses the MLE under the sieve model as a preliminary estimator

¹²For proof, see Theorem 1 in Low (1997).

¹³When g is unknown, one needs to use empirical higher order U -statistic $\hat{\psi}_{m,k}^{\text{emp}}$ to further reduce the bias due to estimating $\Omega_k = \int \bar{z}_k(x)\bar{z}_k(x)g(x)dx$; see Section 1 and Mukherjee, Newey and Robins (2017).

of $\tilde{\theta}_{KL,k}$, then the only difference between MvdV's interval and KBW's universal interval is that for the former the unconditional MLE of $\tilde{\theta}_{KL,k}$ in the numerator is computed from the same sample as the maximum profile likelihood in the denominator, while in the latter they come from different (split) samples. By examining the proofs in MvdV, one can see that the better rate depends crucially on an asymptotic expansion that exploits the fact that the numerator and denominator come from the same sample.

It is yet unclear to us how to reduce the length of a universal confidence interval. On the other hand, we have conjectured in LMR that a nonasymptotic HOIF confidence interval could be constructed by inverting exponential inequalities for U -statistics (Giné, Latała and Zinn, 2000, Adamczak, 2006) but the theory is very challenging and doing so will necessarily increase the confidence interval's length. It will be interesting to investigate if such a nonasymptotic HOIF confidence interval will still shrink faster than the universal confidence interval.

REMARK 5.3 (Final remark on regression and machine learning). We agree with KBW that machine learning is more than prediction. In our paper, we equate “machine learning” with statistical prediction in order to connect with the most current use in causal inference. We believe and expect that many other aspects of machine learning, including clustering, density estimation with generative adversarial networks (GAN), dimension reduction/manifold learning, and optimal transport, will play more and more important roles in causal inference.

ACKNOWLEDGEMENT

Lin Liu and James M. Robins were supported by the U.S. Office of Naval Research Grant N000141912446, and National Institutes of Health (NIH) awards R01 AG057869 and R01 AI127271. Rajarshi Mukherjee's research was partially supported by NSF Grant EAGER-1941419.

SUPPLEMENTARY MATERIAL

Supplement to “Rejoinder: On nearly assumption-free tests of nominal confidence interval coverage for causal parameters estimated by machine learning” (DOI: [10.1214/20-STSS804SUPP](https://doi.org/10.1214/20-STSS804SUPP); .pdf). The online supplementary materials (Liu, Mukherjee and Robins, 2020b) include more technical details for this rejoinder.

REFERENCES

- ADAMCZAK, R. (2006). Moment inequalities for U -statistics. *Ann. Probab.* **34** 2288–2314. [MR2294982](https://doi.org/10.1214/009117906000000476) <https://doi.org/10.1214/009117906000000476>
- ALI, S. M. and SILVEY, S. D. (1966). A general class of coefficients of divergence of one distribution from another. *J. Roy. Statist. Soc. Ser. B* **28** 131–142. [MR0196777](https://doi.org/10.2307/2343131)
- BARRON, A. R. and KLUSOWSKI, J. M. (2018). Approximation and estimation for high-dimensional deep learning networks. Preprint. Available at [arXiv:1809.03090](https://arxiv.org/abs/1809.03090).
- BELLONI, A., CHERNOZHUKOV, V., CHETVERIKOV, D. and KATO, K. (2015). Some new asymptotic theory for least squares series: Pointwise and uniform results. *J. Econometrics* **186** 345–366. [MR3343791](https://doi.org/10.1016/j.jeconom.2015.02.014) <https://doi.org/10.1016/j.jeconom.2015.02.014>
- CAI, T. T., LEVINE, M. and WANG, L. (2009). Variance function estimation in multivariate nonparametric regression with fixed design. *J. Multivariate Anal.* **100** 126–136. [MR2460482](https://doi.org/10.1016/j.jmva.2008.03.007) <https://doi.org/10.1016/j.jmva.2008.03.007>
- CHERNOZHUKOV, V., NEWWEY, W. K. and SINGH, R. (2018). Learning L2 Continuous Regression Functionals via Regularized Riesz Representers. Preprint. Available at [arXiv:1809.05224](https://arxiv.org/abs/1809.05224).
- CSISZÁR, I. (1963). Eine informationstheoretische Ungleichung und ihre Anwendung auf den Beweis der Ergodizität von Markoffschen Ketten. *Magy. Tud. Akad. Mat. Kut. Intéz. Közl.* **8** 85–108. [MR0164374](https://doi.org/10.1016/j.jmva.2008.03.007)
- CUI, Y. and TCHETGEN TCHETGEN, E. (2019). Selective machine learning of doubly robust functionals. Preprint. Available at [arXiv:1911.02029](https://arxiv.org/abs/1911.02029).
- GINÉ, E., LATAŁA, R. and ZINN, J. (2000). Exponential and moment inequalities for U -statistics. In *High Dimensional Probability, II* (Seattle, WA, 1999). *Progress in Probability* **47** 13–38. Birkhäuser, Boston, MA. [MR1857312](https://doi.org/10.1007/978-1-4612-0384-9_2)
- HAYAKAWA, S. and SUZUKI, T. (2020). On the minimax optimality and superiority of deep neural network learning over sparse parameter spaces. *Neural Netw.* **123** 343–361. <https://doi.org/10.1016/j.neunet.2019.12.014>
- KENNEDY, E. H., BALAKRISHNAN, S. and WASSERMAN, L. (2020). Discussion of “On nearly assumption-free tests of nominal confidence interval coverage for causal parameters estimated by machine learning” by L. Liu, R. Mukherjee, and J. Robins. *Statist. Sci.* **35** 540–544.
- LIU, L., MUKHERJEE, R. and ROBINS, J. M. (2020a). An assumption-lean skepticism test of inference validity for doubly robust functionals. Technical report. Available upon request.
- LIU, L., MUKHERJEE, R. and ROBINS, J. M. (2020). Supplement to “Rejoinder: On nearly assumption-free tests of nominal confidence interval coverage for causal parameters estimated by machine learning.” <https://doi.org/10.1214/20-STSS804SUPP>
- LIU, L., MUKHERJEE, R. and ROBINS, J. M. (2020). On nearly assumption-free tests of nominal confidence interval coverage for causal parameters estimated by machine learning. *Statist. Sci.* **35** 518–539.
- LOW, M. G. (1997). On nonparametric confidence intervals. *Ann. Statist.* **25** 2547–2554. [MR1604412](https://doi.org/10.1214/aos/1030741084) <https://doi.org/10.1214/aos/1030741084>
- MUKHERJEE, R., NEWWEY, W. K. and ROBINS, J. M. (2017). Semiparametric efficient empirical higher order influence function estimators. Preprint. Available at [arXiv:1705.07577](https://arxiv.org/abs/1705.07577).
- MURPHY, S. A. and VAN DER VAART, A. W. (2000). On profile likelihood. *J. Amer. Statist. Assoc.* **95** 449–485. [MR1803168](https://doi.org/10.2307/2669386) <https://doi.org/10.2307/2669386>
- NEWWEY, W. K. and ROBINS, J. M. (2018). Cross-fitting and fast remainder rates for semiparametric estimation. Preprint. Available at [arXiv:1801.09138](https://arxiv.org/abs/1801.09138).
- RINALDO, A., WASSERMAN, L. and G'SELL, M. (2019). Bootstrapping and sample splitting for high-dimensional, assumption-lean inference. *Ann. Statist.* **47** 3438–3469. [MR4025748](https://doi.org/10.1214/18-AOS1784) <https://doi.org/10.1214/18-AOS1784>
- ROBINS, J. M. (1988). Confidence intervals for causal parameters. *Stat. Med.* **7** 773–785.

- ROBINS, J., LI, L., TCHETGEN TCHETGEN, E. and VAN DER VAART, A. (2008). Higher order influence functions and minimax estimation of nonlinear functionals. In *Probability and Statistics: Essays in Honor of David A. Freedman* 335–421. IMS, Beachwood, OH, USA.
- ROBINS, J., TCHETGEN TCHETGEN, E., LI, L. and VAN DER VAART, A. (2009). Semiparametric minimax rates. *Electron. J. Stat.* **3** 1305–1321. MR2566189 <https://doi.org/10.1214/09-EJS479>
- ROBINS, J., LI, L., TCHETGEN TCHETGEN, E. and VAN DER VAART, A. (2016). Technical Report: Higher Order Influence Functions and Minimax Estimation of Nonlinear Functionals. Preprint. Available at [arXiv:1601.05820](https://arxiv.org/abs/1601.05820).
- ROBINS, J. M., LI, L., MUKHERJEE, R., TCHETGEN, E. T. and VAN DER VAART, A. (2017). Minimax estimation of a functional on a structured high-dimensional model. *Ann. Statist.* **45** 1951–1987. MR3718158 <https://doi.org/10.1214/16-AOS1515>
- ROTNITZKY, A., SMUCLER, E. and ROBINS, J. M. (2019). Characterization of parameters with a mixed bias property. Preprint. Available at [arXiv:1904.03725](https://arxiv.org/abs/1904.03725).
- SCHMIDT-HIEBER, J. (2020). Nonparametric regression using deep neural networks with ReLU activation function. *Ann. Statist.*. To appear.
- SHEN, Y., GAO, C., WITTEN, D. and HAN, F. (2019). Optimal estimation of variance in nonparametric regression with random design. Preprint. Available at [arXiv:1902.10822](https://arxiv.org/abs/1902.10822).
- WANG, L., BROWN, L. D., CAI, T. T. and LEVINE, M. (2008). Effect of mean on variance function estimation in nonparametric regression. *Ann. Statist.* **36** 646–664. MR2396810 <https://doi.org/10.1214/009053607000000901>
- WASSERMAN, L., RAMDAS, A. and BALAKRISHNAN, S. (2020). Universal inference. *Proc. Natl. Acad. Sci. USA* **117** 16880–16890. <https://doi.org/10.1073/pnas.1922664117>