Published in final edited form as:

J Am Stat Assoc. 2018 ; 113(522): 780–788. doi:10.1080/01621459.2017.1285777.

Parametric-rate inference for one-sided differentiable parameters

Alexander R. Luedtke^{*} and Mark J. van der Laan[†]

Division of Biostatistics, University of California, Berkeley

Abstract

Suppose one has a collection of parameters indexed by a (possibly infinite dimensional) set. Given data generated from some distribution, the objective is to estimate the maximal parameter in this collection evaluated at the distribution that generated the data. This estimation problem is typically non-regular when the maximizing parameter is non-unique, and as a result standard asymptotic techniques generally fail in this case. We present a technique for developing parametric-rate confidence intervals for the quantity of interest in these non-regular settings. We show that our estimator is asymptotically efficient when the maximizing parameter is unique so that regular estimation is possible. We apply our technique to a recent example from the literature in which one wishes to report the maximal absolute correlation between a prespecified outcome and one of p predictors. The simplicity of our technique enables an analysis of the previously open case where p grows with sample size. Specifically, we only require that $\log p$ grows slower than $\sqrt{n}$, where n is the sample size. We show that, unlike earlier approaches, our method scales to massive data sets: the point estimate and confidence intervals can be constructed in $O(np)$ time.

Keywords

stabilized one-step estimator; non-regular inference; variable screening

1 Introduction

Many semiparametric and nonparametric estimation problems yield estimators which achieve a parametric rate of convergence. These estimators are often asymptotically linear, in that they can be written as an empirical mean of an influence function applied to the data. Valid choices of the influence function can be derived as gradients for a functional derivative of the parameter of interest. Applying the central limit theorem then immediately yields Wald-type confidence intervals which achieve the desired parametric rate. Such problems have been studied in depth over the past several decades [Pfanzagl, 1990, van der Vaart, 1991, Bickel et al., 1993, van der Laan and Robins, 2003].

While remarkably general, these approaches rely on the key condition that the parameter of interest is sufficiently differentiable at the data generating distribution for such a gradient to exist. Statisticians are increasingly encountering problems for which parametric-rate estimation is theoretically possible but the parameter is insufficiently differentiable to yield a

^{*}This author gratefully acknowledges the support of the Berkeley Fellowship.

[†]This author gratefully acknowledges the support of NIH grant R01 AI074345-06.

standard first-order expansion demanded by older techniques. For example, suppose we observe baseline covariates, a binary treatment, and an outcome occurring after treatment. We wish to learn the mean outcome under the optimal individualized treatment strategy, i.e. the treatment strategy which makes treatment decisions which are allowed to use baseline covariate information to make treatment decisions Chakraborty and Moodie [2013]. As another example, suppose we observe a vector of covariates $(X_1, \dots, X_p)$ and an outcome Y . We wish give a confidence interval the maximal absolute correlation between a covariate X_k and Y . The lower bound for this quantity is of particular interest since this will suffice for a variable screening procedure. Alternatively, we may only wish to test the null hypothesis that the maximal absolute correlation is zero. McKeague and Qian [2015] provide a test of this null hypothesis using an adaptive resampling test (ART).

These problems belong to a larger class of problems in which one observes $O_1, \dots, O_n$ drawn independently from a P_0 in some (possibly nonparametric) statistical model $\mathcal{M}$ and wishes to estimate

$$\Psi_n(P) \equiv \max_{d \in \mathcal{D}_n} \Psi^d(P), \quad (1)$$

at $P = P_0$, where $\mathcal{D}_n$ is an index set that may rely on sample size and each $\Psi^d: \mathcal{M} \rightarrow \mathbb{R}$ is a sufficiently differentiable parameter to permit parametric-rate estimation using classical methods such as those presented in Bickel et al. [1993]. When there is no unique maximizer $d \in \mathcal{D}_n$ of $\Psi^d(P)$, then the inference problem is typically non-regular, in the sense that the parameter $P \mapsto \max_{d \in \mathcal{D}_n} \Psi^d(P)$ is not sufficiently differentiable to allow the use of standard influence function based techniques for obtaining inference. Hirano and Porter [2012] showed that regular and asymptotically linear estimators fail to exist for such *one-sided pathwise differentiable* parameters. In univariate calculus, functions such as $f(x) = \max\{x, 0\}$ are one-sided differentiable at zero in that the left and right limits of $[f(x + \varepsilon) - f(x)]/\varepsilon$ are well-defined but disagree. The same holds for the Ψ_n evaluated at a distribution P_0 , but now the one-sided differentiability is caused by the subset of $\mathcal{D}_n$ containing the indices which maximize the expression on the right in (1). A small fluctuation in P_0 can greatly reduce the subset of maximizing indices, leading to different derivatives depending on the fluctuation taken.

In this work, we present a method which, loosely, splits the sample in such a way that the estimated index in $\mathcal{D}_n$ which maximizes $\Psi^d(P_0)$ is conditioned on so that this estimated index need not have a limit. We do this iteratively to ensure that our estimator gets the full benefit of the sample size n . When the parameter is fixed with sample size and the d maximizing $\Psi^d(P_0)$ is fixed, we show that our estimator is asymptotically efficient, and therefore also regular. When the maximizing index is not unique, our estimator will not typically be regular. Thus our estimator adapts to the non-regularity of the estimation problem.

Our estimator is inspired by the online estimator for pathwise differentiable parameters presented in van der Laan and Lendle [2014] and a subsequent modification of this estimator in Luedtke and van der Laan [2016] to deal with the non-regularity when estimating the mean outcome under an optimal treatment rule. Such estimators are designed to be efficient in both computational complexity and storage requirements. We show that the estimator that we present in this work inherits many of these computational efficiency properties. We apply our technique to estimate the maximal absolute correlation considered in McKeague and Qian [2015]. In this problem, we show that our estimator runs efficiently in both dimension and sample size, with a runtime of $O(np)$. In practice, this means that the lead author can implement our estimator using only R code and screen $p = 100\,000$ variables using $n = 1\,000$ samples on a single core of his laptop in under a minute. Thus our estimator seems to have both the statistical efficiency that has been demanded of estimators for generations and the computational efficiency that is becoming increasingly important in this new big data era. While the method of McKeague and Qian [2015] can also be implemented in $O(np)$ time when the number of (double) bootstrap samples remains constant, the implicit constant in this procedure implies a much longer runtime.

2 Toy Example

We first present a toy example that we will use to facilitate the presentation of our estimator. While this toy example is simple enough that one can fairly easily come up with alternative estimation strategies, we believe it provides a useful starting point for presenting our general estimation scheme. Suppose we observe an i.i.d. sample $\{O_j \equiv (O_{j,1}, O_{j,2}) : j = 1, \dots, n\}$ of $\mathbb{R}^2$ -valued observations, drawn from some distribution P_0 with $\mathbb{R}^2$ -valued mean $(\Psi_1(P_0), \Psi_2(P_0)) \equiv E_{P_0}[O]$. For general P , similarly define $(\Psi_1(P), \Psi_2(P)) \equiv E_P[O]$. Our objective is to estimate $\Psi(P_0) \equiv \max_{d \in \{1, 2\}} \Psi_d(P_0)$. If $\Psi_1(P_0) = \Psi_2(P_0)$, then this parameter is one-sided but not two-sided pathwise differentiable at P_0 and no regular and asymptotically linear estimator exists [Hirano and Porter, 2012].

To give a reader a sense of the challenges faced by the most intuitive estimation strategy, consider the plug-in estimator that estimates $\Psi_1(P_0)$ and $\Psi_2(P_0)$ using the empirical means of the observations, which we denote $\hat{\mu}_1$ and $\hat{\mu}_2$, and then the estimates $\hat{\psi}_n = \max\{\hat{\mu}_1, \hat{\mu}_2\}$. We have that

$$n^{1/2}[\hat{\psi}_n - \Psi(P_0)] = n^{1/2}[\hat{\mu}_2 - \Psi_2(P_0)] + I\{n^{1/2}[\hat{\mu}_1 - \hat{\mu}_2] \geq 0\} n^{1/2}[\hat{\mu}_1 - \hat{\mu}_2 - \Psi_1(P_0) + \Psi_2(P_0)].$$

By the central limit theorem, $n^{1/2}[\hat{\mu}_d - \Psi_d(P_0) : d = 1, 2]$ converges to a multivariate normal $Z = (Z_1, Z_2)$ with estimable covariance matrix. If $\Psi_1(P_0) > \Psi_2(P_0)$ or $\Psi_2(P_0) > \Psi_1(P_0)$, then the indicator above converges in probability to one or zero, respectively, and the above converges in distribution to Z_1 or Z_2 , respectively. Both of these normal limits can be consistently estimated from the data, and if one knows that $\Psi_1(P_0) > \Psi_2(P_0)$ then the correct index for the limit is consistently estimated by $I\{\hat{\mu}_1 > \hat{\mu}_2\}$. Often one may not be willing to assume that $\Psi_1(P_0) > \Psi_2(P_0)$. To see the challenge that arises, note that if $\Psi_1(P_0) = \Psi_2(P_0)$,

then the right-hand side above converges to $\max\{Z_1, Z_2\}$, which is a non-normal limit. This gives some intuition on the non-regularity of $\Psi(P_0)$ when $\Psi_1(P_0) = \Psi_2(P_0)$: an arbitrarily small shift in $\Psi_1(P_0)$ or $\Psi_2(P_0)$ dramatically changes the limiting behavior of the estimator. Furthermore, it is not in general clear which limiting result one should use to approximate the distribution of $\tilde{\psi}_n$ if the $\Psi_1(P_0) = \Psi_2(P_0) + \varepsilon$ for some small $\varepsilon > 0$, so that asymptotically the limit of the estimator is Z_1 but in practice $\max\{Z_1, Z_2\}$ better approximates the variability of the estimator.

To avoid this problem, we develop an estimator that naturally adapts to the (non-)regularity of the problem. For each $2 \leq j \leq n$, let $(\mu_{j,1}, \mu_{j,2})$ represent the empirical mean of $O_1, \dots, O_j$ and $d_j \equiv \arg\max_d \mu_{j,d}$. Our estimator takes the form $\psi_n \equiv \sum_{j=2}^{n-1} w_j O_{j+1, d_j}$ for positive convex weights $w_2, \dots, w_{n-1}$ that we now define. To expedite the presentation of this toy example, suppose that we know the variance of the first and second components of $O \sim P_0$. Denote these variances by Σ_1^2 and Σ_2^2 , where we assume that $\Sigma_1, \Sigma_2 \in (0, \infty)$. Let $\bar{\Sigma}_n$ denote the harmonic mean of $\Sigma_{d_2}, \dots, \Sigma_{d_{n-1}}$, i.e. $\bar{\Sigma}_n \equiv \left(\frac{1}{n-2} \sum_{j=2}^{n-1} \Sigma_{d_j}^{-1} \right)^{-1}$. Our convex weights are given by $w_j = \bar{\Sigma}_n \Sigma_{d_j}^{-1}$.

We now aim to understand the variability of our estimator. Note that

$$\psi_n - \Psi(P_0) = \sum_{j=2}^{n-1} w_j \left[O_{j+1, d_j} - \Psi_{d_j}(P_0) \right] + \sum_{j=2}^{n-1} w_j \left[\Psi_{d_j}(P_0) - \Psi(P_0) \right].$$

Our analysis will use that the weights w_j are finite because $\Sigma_1, \Sigma_2 \in (0, \infty)$. We first consider the second term on the right. If $\Psi_1(P_0) = \Psi_2(P_0)$, then this term is exactly zero. Suppose $\Psi_1(P_0) \neq \Psi_2(P_0)$. Noting that $\mu_{j,d} \rightarrow \Psi_d(P_0)$ almost surely as $j \rightarrow \infty$, it then follows that, with probability 1, $d_j = \arg\max_d \Psi_d(P_0)$ for all j large enough. It readily follows that $n^{1/2}$ times the second term on the right converges to zero almost surely, and therefore also in probability. Multiplying this term by the random but finite quantity $\bar{\Sigma}_n^{-1}$ does not change this convergence to zero. Thus we have shown that

$$n^{1/2} \bar{\Sigma}_n^{-1} [\psi_n - \Psi(P_0)] = n^{1/2} \bar{\Sigma}_n^{-1} \sum_{j=2}^{n-1} w_j \left[O_{j+1, d_j} - \Psi_{d_j}(P_0) \right] + o_P(1).$$

Finally, we note that the first term above is a martingale sum, and that each term in this sum has variance 1 thanks to the choice of weights. A standard martingale central limit theorem [e.g., Gaenssler et al., 1978] then shows that this term converges to a standard normal random variable, and standard arguments show that the interval $[\psi_n \pm 1.96 \bar{\Sigma}_n^{-1/2}]$ contains $\Psi(P_0)$ with probability approaching 0.95.

3 Estimator

We will now present our technique for a general estimation problem. We now introduce the notion of pathwise differentiability, since this provides the key object needed to construct our estimator.

3.1 Pathwise differentiability

We assume that each parameter Ψ^d , $d \in \mathcal{D}_n$ for any n , is pathwise differentiable for all distributions in our model [see, e.g., Pfanzagl, 1990, Bickel et al., 1993]. For each $P \in \mathcal{M}$, we let $D^d(P)$ denote the canonical gradient of Ψ^d at P . By definition $D^d(P)(O)$ is mean zero with finite variance under sampling from P . Typically pathwise differentiability implies that Ψ^d satisfies the following linear expansion for any $P \in \mathcal{M}$ and $d \in \mathcal{D}_n$:

$$\Psi^d(P) - \Psi^d(P_0) = - \int D^d(P)(o) dP_0(o) + \text{Rem}_n^d(P), \quad (2)$$

where we omit the dependence of $\text{Rem}_n^d(P)$ on P_0 in the notation and indicate its possible dependence on sample size with the subscript n . Above $\text{Rem}_n^d(P)$ is a second-order remainder term that is small whenever P is close to P_0 . We consider this condition more closely in our example, but for non-sample size dependent parameters this term can typically be made to be $O_P(1/n)$ in a parametric model and often can be made to be $o_P(1/\sqrt{n})$ in a nonparametric model. In the toy example from Section 2, the two parameters $\Psi_d(P_0)$ are linear, and so $D^d(P)(o) = o_d - E_P[O_d]$ and $\text{Rem}_n^d(P) = 0$ for all d, P . For a more thorough presentation, see Pfanzagl [1990] or Bickel et al. [1993].

Knowing the canonical gradient of a parameter enables one to implement a one-step estimator [see Section 5.7 of van der Vaart, 1998]. To ease discussion, fix d . Suppose one has an initial estimate $\hat{P}_n$ of the components of P_0 needed to evaluate Ψ^d and D^d . Then a one-step estimate of $\Psi^d(P_0)$ is given by $\tilde{\Psi}_n^d = \Psi^d(\hat{P}_n) + \frac{1}{n} \sum_{i=1}^n D^d(\hat{P}_n)(O_i)$. Under empirical process and consistency conditions on $\hat{P}_n$, one can show that $n^{1/2}[\tilde{\Psi}_n^d - \Psi^d(\hat{P}_n)]$ converges in distribution to a normal random variable with estimable variance. In the next section, we present a variant of this estimator that allows for the selection of the optimizing d , even when the optimal index is non-unique.

3.2 Estimator and confidence interval

We now present a stabilized one-step estimator for problems of the type found in (1) when the required differentiability condition on Ψ^d holds.

Let $\{\ell_n\}$ be some sequence such that $n - \ell_n \rightarrow \infty$. One possible choice is $\ell_n = 0$ for all n . For each $j = \ell_n, \dots, n - 1$, let d_{nj} represent an estimate of a maximizer of (1) obtained using observations $(O_i : i = 1, \dots, j)$, $\hat{P}_{nj}$ be an estimate of P_0 obtained using observations $(O_i : i =$

$1, \dots, j)$, and $\hat{P}_{nj} \hat{D}_{nj}$ equal $D^d(P)$ evaluated at $P = \hat{P}_{nj}$ and $d = d_{nj}$. For nonnegative weights $w_{\ell_n}, \dots, w_{n-1}$ that we will define shortly with $\sum_{j=\ell_n}^{n-1} w_j = n - \ell_n$, our stabilized one-step estimate takes the form

$$\psi_n \equiv \frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} w_j \left[\Psi^{d_{nj}}(\hat{P}_{nj}) + \hat{D}_{nj}(O_{j+1}) \right].$$

Our proposed 95% confidence interval has the form

$$[\text{LB}_n, \text{UB}_n] \equiv \left[\psi_n \pm 1.96 \frac{\bar{\sigma}_n}{\sqrt{n - \ell_n}} \right],$$

where we will define $\bar{\sigma}_n$ momentarily and one can replace 1.96 by the desired quantile of the normal distribution to modify the confidence level.

We now define the weights. Let $\hat{\sigma}_{nj}^2$ represent an estimate of the variance of $\hat{D}_{nj}(O)$, $O \sim P_0$, conditional on observations $O_1, \dots, O_j$. This estimate should only rely on those j observations. Often we can let

$$\hat{\sigma}_{nj}^2 \equiv \frac{1}{j} \sum_{i=1}^j \left[\hat{D}_{nj}(O_i) - \frac{1}{j} \sum_{i=1}^j \hat{D}_{nj}(O_i) \right]^2, \quad j = \ell_n, \dots, n-1.$$

The standard deviation type variable in the confidence interval definition is given by

$\bar{\sigma}_n \equiv \left(\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} \right)^{-1}$, and the weights are given by $w_j \equiv \bar{\sigma}_n \hat{\sigma}_{nj}^{-1}$, where we have omitted the possible dependence of the weights on sample size in the notation.

Our estimator is similar to the online one-step estimator developed in van der Laan and Lendle [2014] for streaming data, but it weights each term proportionally to the estimated inverse standard deviation of $\hat{D}_{nj}(O)$ when $O \sim P_0$. Our confidence interval takes a form similar to a Wald-type confidence interval, but replaces the typical standard deviation with $\bar{\sigma}_n$ and has width on the order of $1/\sqrt{n - \ell_n}$ rather than $1/\sqrt{n}$. Note of course that $\ell_n = o(n)$ implies that $1/\sqrt{n - \ell_n} - 1/\sqrt{n}$ converges to zero.

3.3 Validity of confidence interval

We now prove the validity of our confidence interval. Let $\sigma_{nj}^2 \equiv \text{Var}_{P_0}(\hat{D}_{nj}(O) | O_1, \dots, O_j)$. The validity of the lower bound of the confidence interval relies on the following conditions:

C1) There exists some $M < \infty$ such that

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} P_0 \left(\left| \frac{\widehat{D}_{nj}(O)}{\widehat{\sigma}_{nj}} \right| < M \middle| O_0, \dots, O_{j-1} \right) \rightarrow 1 \text{ in probability as } n \rightarrow \infty.$$

C2) $\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \left| \frac{\sigma_{nj}^2}{\widehat{\sigma}_{nj}^2} - 1 \right| \rightarrow 0$ in probability as $n \rightarrow \infty$.

C3) $\frac{1}{\sqrt{n - \ell_n}} \sum_{j=\ell_n}^{n-1} \widehat{\sigma}_{nj}^{-1} \widehat{\text{Rem}}_{nj} \rightarrow 0$ in probability as $n \rightarrow \infty$, where $\widehat{\text{Rem}}_{nj} \equiv \text{Rem}^{d_{nj}}(\widehat{P}_{nj})$.

The validity of the upper bound requires the following additional condition:

C4) $\frac{1}{\sqrt{n - \ell_n}} \sum_{j=\ell_n}^{n-1} \widehat{\sigma}_{nj}^{-1} \left[\Psi^{d_{nj}}(P_0) - \Psi_n(P_0) \right]$ converges to zero in probability as $n \rightarrow \infty$.

We now present our main result.

Theorem 1—(Validity of confidence interval). *If C1), C2), and C3) hold, then*

$$\liminf_{n \rightarrow \infty} \Pr(\Psi_n(P_0) \geq \text{LB}_n) \geq 1 - \alpha/2.$$

If C4) also holds, then

$$\lim_{n \rightarrow \infty} \Pr(\text{LB}_n \leq \Psi_n(P_0) \leq \text{UB}_n) = 1 - \alpha.$$

Proof: The definition ψ_n combined with (2) yield that

$$\begin{aligned} \sqrt{n - \ell_n} \widehat{\sigma}_n^{-1} [\psi_n - \Psi_n(P_0)] &= \frac{1}{\sqrt{n - \ell_n}} \sum_{j=\ell_n}^{n-1} \widehat{\sigma}_n^{-1} \left(\widehat{D}_{nj}(O_{j+1}) - E_{P_0}[\widehat{D}_{nj}(O) | O_1, \dots, O_j] \right) \\ &+ \frac{1}{\sqrt{n - \ell_n}} \sum_{j=\ell_n}^{n-1} \widehat{\sigma}_n^{-1} \left[\Psi^{d_{nj}}(P_0) - \Psi_n(P_0) + \widehat{\text{Rem}}_{nj} \right]. \end{aligned} \quad (3)$$

The second line converges to zero in probability by C3) and C4). By C1), C2), and the martingale central limit theorem for triangular arrays in Gaenssler et al. [1978], (3) converges in distribution to a standard normal random variable. A standard Wald-type confidence interval construction argument shows that the confidence interval has coverage approaching $1 - \alpha$ under C1) through C4).

Now suppose C4) does not hold. By (1), $\sum_{j=\ell_n}^{n-1} \widehat{\sigma}_{nj}^{-1} \left[\Psi^{d_{nj}}(P_0) - \Psi_n(P_0) \right] \leq 0$. The same argument readily shows the validity of the lower bound under only C1), C2), and C3). $\square$

The conditions of the theorem are discussed in Appendix A.1. Appendix A.2 considers the asymptotic efficiency of our estimator when the parameter in (1) does not rely on sample size. High level conditions are provided, and we then argue that these conditions are plausible when the maximizing index in (1) is unique. Appendix A.3 discusses computationally efficient implementations of our general estimator.

4 Maximal correlation example

4.1 Problem formulation

We now present the running example of this work, namely the maximal correlation estimation problem considered by McKeague and Qian [2015]. The observed data structure is $O = (X, Y)$, where $X = (X_k : k = 1, \dots)$ is a $[-1, 1]^\infty$ vector of predictors and Y is an outcome in $[-1, 1]$. For each n , we let $\mathcal{X}_n$ represent a subset of these predictors of size p , where throughout we assume that

$$\beta_n^2 \equiv \frac{\log p}{\sqrt{n}} \rightarrow 0 \text{ as } n \rightarrow \infty. \quad (4)$$

For readability, we omit the dependence of p on n in the notation. Under a distribution P , the maximal absolute correlation of a predictor with Y is given by

$$\Psi_n(P) \equiv \max_{k \in \mathcal{X}_n} |\text{Corr}_P(X_k, Y)|, \quad (5)$$

where $\text{Corr}_P(X_k, Y)$ is the correlation of X_k and Y under P . We wish to develop confidence intervals for $\Psi_n(P_0)$. When a test of $H_0 : \Psi_n(P_0) = 0$ against the complementary alternative is of interest, we also wish to establish the behavior of our test against local alternatives as was done in McKeague and Qian [2015].

In contrast to McKeague and Qian [2015], the procedure that we present in this work:

1. is proven to work when p grows with sample size at any rate satisfying (4);
2. yields confidence intervals for the maximal correlation rather than just a test of the null hypothesis that it is equal to zero;
3. allows a non-null maximizer in (5) to be non-unique.

While McKeague and Qian argued that 3) is unlikely in practice, having two non-null maximizers be approximately equal may still have finite sample implications for their test in some settings.

We now show that this problem fits in our framework. To satisfy the pathwise differentiability condition, we let and, for each $d = (k, m) \in \mathcal{D}_n$,

$$\Psi^d(P) \equiv m \text{Corr}_P(X_k, Y).$$

Note that $\Psi_{\pi}(P)$ now takes the form in (1), where we note that the use of m in the definition of Ψ^d serves to ensure that $\Psi_{\pi}(P_0)$ represents the correlation with the maximal *absolute value*.

4.2 Differentiability condition

Canonical gradients—For each k , let $s_P^2(X_k) \equiv \text{Var}_P(X_k)$, and likewise for $s_P^2(Y)$. For ease of notation we let $s_0^2(X_k) \equiv s_{P_0}^2(X_k)$, and likewise for $s_0^2(Y)$ and $\text{Corr}_0(X_k, Y)$. An application of the delta method shows that Ψ^d has canonical gradient $D^d(P)(o)$ given by

$$m \times \left(\frac{(x_k - E_P[X_k])(y - E_P[Y])}{s_P(X_k)s_P(Y)} - \frac{1}{2}\text{Corr}_P(X_k, Y) \left[\frac{(x_k - E_P[X_k])^2}{s_P^2(X_k)} + \frac{(y - E_P[Y])^2}{s_P^2(Y)} \right] \right).$$

In order to ensure that $D^d(P_0)$ is uniformly bounded for all d , we assume throughout that, for some $\delta \in (0, 1]$,

$$\min\{s_0(Y), s_0(X_1), s_0(X_2), \dots\} > \delta.$$

Second-order remainder—Fix $d = (k, m) \in \mathcal{D}_n$ and $P \in \mathcal{M}$. Let $\tilde{\delta} \in (0, 1]$ be some constant such that both $s_P(X_k)$ and $s_P(Y)$ are larger than $\tilde{\delta}$. Lemma A.3 in Appendix B.1 proves that

$$\begin{aligned} |\text{Rem}^d(P)| &\leq \tilde{\delta}^{-1} (|s_P(X_k)s_P(Y) - s_0(X_k)s_0(Y)| |\text{Corr}_P(X_k, Y)| \\ &\quad + (E_P[X_k] - E_{P_0}[X_k])^2 + (E_P[Y] - E_{P_0}[Y])^2) \\ &\quad + \frac{s_0^2(Y)}{s_P^2(Y)} [s_P(X_k) - s_0(X_k)]^2 + \frac{s_0^2(X_k)}{s_P^2(X_k)} [s_P(Y) - s_0(Y)]^2. \end{aligned} \quad (6)$$

The first term above is small if $s_P(X_k)$, $s_P(Y)$, and $\text{Corr}_P(X_k, Y)$ are close to $s_0(X_k)$, $s_0(Y)$, and $\text{Corr}_0(X_k, Y)$. The middle terms are small if $E_P[X_k]$ and $E_P[Y]$ are close to $E_{P_0}[X_k]$ and $E_{P_0}[Y]$. The final terms are small if $s_P(X_k)$ and $s_P(Y)$ are close to $s_0(X_k)$ and $s_0(Y)$.

Variance of canonical gradients—For any given d , there is no elegant (and informative) expression for $\text{Var}_{P_0}[D^d(P_j)(O)]$. Nonetheless, we show in Lemma A.6 of

Appendix B.1 that our estimates $\hat{\sigma}_{nj}^2$, taken as the sample variance of $D^{dj}(P_j)(O)$ for an index estimate d_{nj} to be defined in the next subsection, concentrate tightly about σ_{nj}^2 with high probability when the sample size is large enough. Thus, in practice, one can actually

check if σ_{nj}^2 is small by looking at $\hat{\sigma}_{nj}^2$. If P_0 is normal, then this variance is equal to $\left[1 - \text{Corr}_{P_0}(X_k, Y)^2\right]^2$, and so is only zero if $\text{Corr}_{P_0}(X_k, Y) = 1$. Though such an elegant expression does not exist for the variance of $D^d(P_0)(O)$ for general distributions, one can still show in general that the variance of $D^d(P_0)$ is equal to zero only if $\text{Corr}_{P_0}(X_k, Y) = 1$. Here we make the slightly stronger assumption that

$$\inf_{n \geq 2} \min_{(k, m) \in \mathcal{K}_n \times \{-1, 1\}} \text{Var}_{P_0}(D^d(P_0)(O)) \geq \gamma > 0. \quad (7)$$

4.3 Our estimator

We will use the estimator presented in Section 3 to estimate $\Psi_n(P_0)$. At each index $j \in \mathcal{J}_n$ we use the empirical distribution P_j of the observations $O_1, \dots, O_j$ to estimate P_0 . We let our optimal index estimate $d_{nj} \equiv (k_{nj}, m_{nj})$, where $k_{nj} \equiv \arg\max_{k \in \mathcal{K}_n} \text{Corr}_{P_j}(X_k, Y)$ and

$m_{nj} \equiv \text{sgn}[\text{Corr}_{P_j}(X_{k_{nj}}, Y)]$. We estimate σ_{nj}^2 with the variance of $\hat{D}_{nj}(O)$ under P_j .

In Appendix B.1, we detail conditions on $\mathcal{J}_n$ which ensure that $\mathcal{J}_n$ does not grow too slowly or quickly. For any $\varepsilon \in (0, 2)$, one possible choice of $\mathcal{J}_n$ that satisfies these conditions is

$$\mathcal{J}_n = \max\left\{(\log \max\{n, p\})^{1+\varepsilon}, n \exp(-\beta_n^{-2+\varepsilon})\right\}. \quad (8)$$

We show that this choice of $\mathcal{J}_n$ ensures C1), C2), and C3) in Appendix B.1. By Theorem 1 this establishes the validity of the lower bound of our confidence interval. We can also show that this lower bound is tight up to a term of the order $n^{-1/4}\beta_n$.

Theorem 2—(Tightness of the lower bound). *For any sequence $t_n \rightarrow \infty$, $\Psi_n(P_0) < \text{LB}_n + t_n n^{-1/4}\beta_n$ with probability approaching 1.*

We note that the choice of $\mathcal{J}_n$ in (8) is only needed if the set of indices $\mathcal{K}_n$ changes with sample size. For fixed $\mathcal{K}_n$, one could take $\mathcal{J}_n$ fixed, e.g. $\mathcal{J}_n = 2$, and still have a valid lower bound provided the estimates of σ_{nj} are truncated from below at some $\tilde{\delta} > 0$ (see Lemma A.1). While choosing $\mathcal{J}_n$ according to (8) is still advisable since this is what will enable us to study the behavior of a hypothesis testing procedure under local alternatives, this invariance to $\mathcal{J}_n$ should at least reassure the user that most choices of $\mathcal{J}_n$ will perform reasonably well. In Luedtke and van der Laan [2016], we evaluated the stabilized one-step estimator on a variety of choices of $\mathcal{J}_n$ and found little sensitivity to this tuning parameter. Nonetheless, we consider the development of a data adaptive selection procedure for choosing an $\mathcal{J}_n$ satisfying C1), C2), and C3) an important area for future work. In parallel to how McKeague and Qian [2015] used the bootstrap to select their tuning parameter, one might consider using the bootstrap to select $\mathcal{J}_n$, though it remains to determine an appropriate criterion for selecting $\mathcal{J}_n$.

Because our ℓ_n -specific lower bound is defined using a normal limiting result rather than the bootstrap, such a selection procedure would avoid the use of a computationally burdensome double bootstrap.

We now consider the validity of the upper bound of our confidence interval, which holds under C4). This condition is trivially valid if $\Psi_n(P_0) = 0$ for all n . Condition C4) is also valid under the following margin condition:

MC) For some sequence $t_n \rightarrow \infty$, there exists a sequence of non-empty subsets $\mathcal{K}_n^* \subseteq \mathcal{K}_n$ such that, for all n ,

$$\begin{aligned} \sup_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)| - \inf_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)| &= o(n^{-1/2}), \\ \inf_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)| &\geq \sup_{k \in \mathcal{K}_n \setminus \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_k, Y)| + t_n n^{-1/4} \beta_n. \end{aligned}$$

If $\mathcal{K}_n^* = \mathcal{K}_n$, then the supremum over $\mathcal{K}_n \setminus \mathcal{K}_n^*$ is taken to be zero.

Theorem 3—(Validity of the upper bound). *If MC) or $\Psi_n(P_0) = 0$ for all n , then C4) holds so that $\text{LB}_n - \Psi_n(P_0) \leq \text{UB}_n$ with probability approaching $1 - \alpha$.*

We outline the techniques used to prove these two results at the end of this subsection. Complete proofs are given in Appendix B.1.

Suppose we wish to test $H_0 : \Psi_n(P_0) = 0$ against $H_1 : \Psi_n(P_0) > 0$. Consider the test that rejects H_0 if $\text{LB}_n > 0$. We wish to explore the behavior of this test under local alternatives where $\Psi_n(P_0)$ converges to zero slower than $n^{-1/4} \beta_n$. Theorem 2 shows that this test has power converging to one under such local alternatives. Furthermore, as the lower bound is valid in general, this test has type I error of at most $\alpha/2$ under the null. This is indeed an exciting result as it enables the study of local alternatives even when dimension grows quickly with sample size. If dimension does not grow with sample size, this shows that we can detect against any alternatives converging to zero slower than $n^{-1/2} \sqrt{\log n}$. We would not be surprised if the $\sqrt{\log n}$ is unnecessary, but rather that it is simply a result of our proof techniques which give high probability bounds on the concentration of our correlation estimates at each sample size. McKeague and Qian [2015] showed that their method is consistent against a class of alternatives converging to zero slower than $n^{-1/2}$ provided the optimal index is unique. Our result does not rely on this uniqueness condition. We emphasize that we only used MC) to establish the validity of the *upper bound* of our confidence interval. Our lower bound, and therefore our ability to reject the null of uniformly zero correlation, is valid even without this margin condition.

Theorem 3 shows that the upper bound of our confidence interval is also valid under a reasonable margin condition. The margin condition states that there may be many non-null approximate maximizers provided their absolute correlations are well-separated from the absolute correlations of the other predictors with Y . By “approximate” we mean that their

absolute correlations all fall within $\alpha(n^{-1/2})$ of one another. If $\mathcal{K}_n$ does not depend on sample size, then this theorem shows that our two-sided confidence interval is always valid.

Sketch of proofs of Theorems 2 and 3: Our proofs of both of these theorems rely on high-probability bounds of the absolute differences between our estimates of $s_{P_j}^2(X_k)$, $s_{P_j}^2(Y)$, $\text{Corr}_{P_j}(X_k, Y)$, and $\hat{\sigma}_{nj}$ and their population counterparts, uniformly over $k \in \mathcal{K}_n$ and j . We show that, with probability at most $1 - 1/n$, all of these absolute differences are upper bounded by constants (with explicit dependence on γ and δ) times $j^{-1/2} \log \max\{n, p\}$.

Condition C1) follows once we show that, with high probability, $s_{P_j}^2(X_k)$ and $s_{P_j}^2(Y)$ are bounded below by $\delta/2$ and $\hat{\sigma}_{nj}^2$ is bounded below by $\gamma/2$ uniformly over $j \in \mathcal{J}_n$ for n large enough. Condition C2) and C3) are easy consequences of our concentration results. The concentration results also yield that

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} \left[\Psi^{d_{nj}}(P_0) - \Psi_n(P_0) \right] = O_{P_0}(n^{-1/4} \beta_n),$$

which then quickly yields Theorem 2 thanks to the expression in (3).

Now suppose MC) holds. By our concentration inequalities, we select a $k_{nj} \in \mathcal{K}_n^*$ for each $j \geq Ct_n^{-1}n$ with high probability, where C is a constant. We also correctly specify m_{nj} to be the sign of $\text{Corr}_0(X_{k_{nj}}, Y)$. Because all of the absolute correlations in $\mathcal{K}_n$ are small, the difference between $\Psi^{d_{nj}}(P_0)$ for $d_{nj} = (k_{nj}, m_{nj})$ and $\Psi_n(P_0)$ is very small. If $\ell_n < Ct_n^{-1}$, then we can apply our concentration inequalities to establish that these first few values of j for which $j < Ct_n^{-1}$ are small enough so that C4) still holds, yielding Theorem 3. $\square$

In Appendix B.2, we show that our estimator runs in $O(np)$ time. We show that the estimate can be computed using $O(p)$ storage when the observations $O_1, \dots, O_n$ arrive in a data stream. This result is closely related to the fact that, for a $\mathbb{R}^p$ -valued sequence $\{t_i\}$, the sum $S_j \equiv \sum_{i=1}^j t_i$ at $j = n$ can be computed in time $O(np)$ using storage $O(p)$. In particular, one can use the recursion relation $S_j = t_j + S_{j-1}$, thereby only storing t_j and S_{j-1} when computing S_j . Our estimate can also be computed in $O(np)$ time and $O(n)$ storage when the vectors $(X_{jr} : j = 1, \dots, n) \in \mathbb{R}^n$ arrive in a stream for $r = 1, 2, \dots, p$, where X_{jr} is the observation of X_r for individual j . We do not prove the $O(n)$ storage result in the appendix due to space constraints, though the algorithm is closely related to that given in Appendix B.2.

5 Simulation study

We now consider the power and scalability of our method using the simulations similar to those described in McKeague and Qian [2015]. Let $X \sim \text{MVN}(0, \Sigma)$ for Σ a $p \times p$ covariance

matrix to be given shortly, and $\tau_1, \dots, \tau_p$ be a sequence of i.i.d. normal random variables independent of all other quantities under consideration. We will use two types of errors: the homoscedastic error τ_1 and the heteroscedastic error $\eta(X) \equiv \sum_{k=1}^p X_k \tau_k / \sqrt{p}$. For $(n, p) = (200, 200), (500, 2000)$, we generate data using the following distributions: (N.IE) $Y = \tau_1$, (A1.IE) $Y = X_1/5 + \tau_1$, (A2.IE) $Y = 0.15 \sum_{k=1}^5 X_k - 0.1 \sum_{k=6}^{10} X_k - \tau_1$, (N.DE) $Y = \eta(X)$, (A1.DE) $Y = X_1/5 + \eta(X)$, and (A2.DE) $Y = 0.15 \sum_{k=1}^5 X_k - 0.1 \sum_{k=6}^{10} X_k + \eta(X)$. For $(n, p) = (2\,000, 30\,000)$, we generate data using the following distributions: (N.IE) $Y = \tau_1$, (A3.IE) $Y = X_1/15 + \tau_1$, and $Y = 0.03 \sum_{k=1}^5 X_k - 0.015 \sum_{k=6}^{10} X_k + \tau_1$. We set all of the diagonal elements in the covariance matrix Σ equal 1, and the off-diagonal elements equal ρ , where for each simulation setting we let $\rho = 0, 0.25, 0.5, 0.75$. Unless otherwise specified, all simulations are run using 1 000 Monte Carlo simulations in R [R Core Team, 2014]. Code is available in the Supplementary Materials.

We conduct a 5% test of $\Psi(P_0) > 0$ by checking if the lower bound of a 90% confidence interval for this quantity is greater than zero. We use models N.IE and N.DE to evaluate type I error and all other models evaluate power. We run our method with ℓ_n as in (8), where we let $\varepsilon = 0.5$. For ease of implementation, we compute our method on chunks of data of size $(n - \ell_n)/10$ (see Section 6.1 of Luedtke and van der Laan, 2016).

We compare our method to the ART of McKeague and Qian [2015]. The ART relies on a tuning parameter λ_n satisfying $\lambda_n/\sqrt{n} \rightarrow 0$ and $\lambda_n \rightarrow \infty$ that is selected via a double bootstrap procedure. We implemented code that we obtained from the authors (McKeague and Qian) that selects $\lambda_n = \sqrt{a \log n}$ from a grid of a varying between 0.5 and 4. Due to computational limitations, we ran 400 outer bootstrap samples and 200 inner bootstrap samples (rather than the default of 1 000 samples for both layers of bootstrap), and also reduced the grid for a from the default (0.5, 0.55, ..., 4) to (0.5, 0.6, ..., 4). We also reduced the number of Monte Carlo replicates for the ART to 200 and only ran ART on the smallest sample size $(n, p) = (200, 200)$. While we were not able to run the double bootstrap at the moderate sample size $(n, p) = (500, 2000)$ due to computational constraints, we were able to mimic the double bootstrap procedure by selecting an oracle choice of $\lambda_n = \sqrt{a \log n}$. In particular, we ran ART for the fixed choices of $a = 0.5, 2.25, 4$, found that $a = 4$ appropriately controlled type I error while the other choices of a typically did not, and reported the results of ART at this fixed tuning parameter. We were unable to run even the oracle procedure at the largest sample size with due to computational constraints.

We also compared our procedure to the analogue of ART described in Section 2 of Zhang and Laber [2015], where this analogue does not require running a double bootstrap. This latter procedure is referred to as the “parametric bootstrap” in Zhang and Laber [2015], though to avoid confusion with other bootstrap procedures here we refer to their method as “ZL”. The ZL procedure assumes a locally linear model with homoscedastic errors. Note that the homoscedasticity requirement is stronger than the uncorrelated error requirement made by the ART. In fact, the errors are guaranteed to be uncorrelated with the predictors under the null of zero maximal absolute correlation, thereby ensuring the type I error control

of ART. We use 500 bootstrap draws for each run of the ZL procedure. Zhang and Laber show that their method, which does not involve running a computationally burdensome double bootstrap procedure, has comparable performance to ART across sample sizes and predictor dimension, while being more computationally efficient. The ZL procedure is less computationally intensive than the ART, but still requires estimating the $p \times p$ covariance matrix Σ and simulating from a $N(0, \widehat{\Sigma})$ distribution. Due to computational constraints, we only run ZL for $p = 2\,000$ and not for $p = 30\,000$. We also compare our method to a Bonferroni-corrected t -test.

Figure 1 displays the power of the four testing procedures for (n, p) equal to $(200, 200)$ and $(500, 2000)$ for the homoscedastic data generating distributions N.IE, A1.IE, and A2.IE. The ART and ZL procedures perform best in both of these settings. We can show (details omitted) that our method underperforms in this setting due to the second-order term representing the cost for estimating d_0 on subsets of the data of size $j \ll n$ early on in the procedure. While Theorem A.11 ensures that the estimate of d_0 will be asymptotically valid, there appears to be a noticeable price to pay at small sample sizes.

Figure 2 displays the power of the three testing procedures for (n, p) equal to $(200, 200)$ and $(500, 2000)$ for the heteroscedastic data generating distributions. The ZL procedure fails to control the type I error in this setting. This is unsurprising given that this test was developed under a local linear model with independent errors. All other methods adequately control type I error in this setting, especially at the larger sample size $n = 500$, while we see that the Bonferroni and ART procedures achieve slightly better power than our method for these data generating distributions.

Figure 3 displays the power of our method and the Bonferroni procedure for (n, p) equal to $(2\,000, 30\,000)$. While (unsurprisingly) Bonferroni performs well when the correlation between the predictors in X is low, our method outperforms the Bonferroni procedure when the correlation increases. We expect that, were we able to run ART or ZL at this sample size, they would outperform all other methods under consideration as they did at the smaller sample sizes. Nonetheless, both methods quickly become computationally impractical when p gets large, whereas our procedure and the Bonferroni procedure can still be implemented at these sample sizes.

We also ran our method at different choices of ℓ_n for $(n, p) = (200, 200)$ and $(500, 2\,000)$ (details not shown), namely defined according to (8) with $\varepsilon = 0.25, 1, 1.5, 1.75$. We found little sensitivity to the choice of ε , with the exception that choosing $\varepsilon = 1.75$ often led to a moderate loss of power (at most 15% on an additive scale). This is not surprising given that, at $\varepsilon = 1.75$, ℓ_n is approximately equal to $n/2$ for both (n, p) settings.

6 Discussion

We have presented a general method for estimating the (possibly non-unique) maximum of a family of parameter values indexed by $d \in \mathcal{D}_n$. Such an estimation problem is generally non-regular because minor fluctuations of the data generating distribution can change the subset of $\mathcal{D}_n$ for which the corresponding parameter is maximized. Our estimate takes the form of a

sum of the terms of a martingale difference sequence, which quickly allows us to apply the relevant central limit theorem to study its asymptotics and develop Wald-type confidence intervals. The estimator adapts to the non-regularity of the problem, in the sense that we can give reasonable conditions under which it is regular and asymptotically linear when the maximizer is unique so that regularity is possible.

We have applied our approach to the example of McKeague and Qian [2015] in which one wishes to learn about the maximal absolute correlation between a prespecified outcome and a predictor belonging to some set. The sample splitting that is built into our estimator has enabled us to analyze the estimator when the dimension p of the predictor grows with sample size slowly enough so that $n^{-1/2} \log p \rightarrow 0$ as n goes to infinity. While McKeague and Qian focus on testing the null hypothesis that this maximal absolute correlation is zero, we have established valid confidence intervals for this quantity. The lower bound of our confidence interval is particularly interesting because it is valid under minimal conditions. When p is very large, one might expect that the null of no correlation between the outcome and any of the predictors is unlikely to be true. In these problems, having an estimate of the maximal absolute correlation, or at least a lower bound for this quantity, will likely still be interesting as a measure of the overall relationship between X and Y .

We have also studied the behavior of this null hypothesis test under local alternatives, showing that our test is consistent when the maximal absolute correlation shrinks to zero slower than $n^{-1/2}(\log \max\{n, p\})^{1/2}$. When the dimension of the predictor is fixed, the test of McKeague and Qian is consistent against alternatives shrinking to zero more slowly than $n^{-1/2}$ rather than $(\log n)^{1/2}n^{-1/2}$. We would not be surprised to find that this $(\log n)^{1/2}$ is unnecessary for p fixed and can be removed using more refined proof techniques.

McKeague and Qian do not require that Y and the coordinates of X have range in $[-1, 1]$. We have made this boundedness assumption out of convenience for our proofs and expect that we can replace the boundedness assumptions with appropriate moment assumptions without significantly changing the results. Our simulation results support this claim. The boundedness condition is not as restrictive as it may first seem, as unbounded X and Y can be rescaled to be bounded. Since the sharp null $H_0 : \Psi_{\pi}(P_0) = 0$ is invariant to strictly monotonic transformations of X and Y , our theoretical results yield a valid H_0 test after applying, e.g., the sigmoid transformation to X and Y .

We note that, in our simulations, ART and ZL achieve the highest power among competing methods, though for our heteroscedastic simulation setting ZL failed to control the type I error. We were not able to run either of these methods at our largest sample size due to computational constraints. The ZL procedure as currently described is computationally expensive and does not scale well to large data sets, especially when the dimension of the predictor p is large. This difficulty occurs because the procedure requires the computation of a $p \times p$ covariance matrix. The ART method presented in McKeague and Qian [2015], which achieves similar power to ZL, is in practice even more computationally burdensome due to its use of a double bootstrap. Nonetheless, from a theoretical computational complexity standpoint, the ART method can be made to scale as $O(np)$ provided the number of bootstrap draws remains fixed. Though the number of bootstrap samples will likely be fixed in

practical applications, we note that ART cannot maintain consistency against local alternatives unless the number of bootstrap samples grows with sample size, thereby yielding a slower than order- np runtime. As is to be expected from marginal screening procedures that perform an $O(n)$ screening operation p times, our method attains an $O(np)$ runtime. This computational efficiency, combined with the asymptotic theory supporting our method's power against local alternatives under increasing covariate dimension and efficiency under fixed alternatives and covariate dimension, demonstrates what is achievable by marginal screening procedures. Given our simulations, we also believe that developing rigorous asymptotic theory under increasing dimension for the ART methods is an important area for future work.

The stabilized one-step estimator presented in this paper applies to many other situations not considered in this paper. In an earlier work, we showed that this estimator is useful for estimating the mean outcome under an optimal individualized treatment strategy Luedtke and van der Laan [2016], where the class $\mathcal{D}_n$ now indexes functions mapping from the covariate space to the set of possible treatment decisions. Thanks to the martingale structure of our estimator, the stabilized one-step estimator can be used to construct confidence intervals when the data is drawn sequentially so that the data generating distribution for observation j can depend on that of the first $j - 1$ observations. One interesting example along these lines is to obtain inference for the value of the optimal arm in a multi-armed bandit problem, even in the case where the optimal arm is non-unique and the reward distributions for the optimal arms have different variances. We look forward to seeing further applications of the general template for a stabilized one-step estimator that we have presented in this paper.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Appendix A General estimator

A.1 Discussion of conditions of Theorem 1

In this section, we consider the setting where the parameter in (1) does not depend on sample size, and consequently omit the n subscript to quantities which no longer depend on sample size. We will show that C7) and the following conditions imply the conditions of Theorem 1:

$$\text{C9)} \quad \hat{\sigma}_j^2 - \sigma_j^2 \text{ converges to zero in probability as } j \rightarrow \infty.$$

$$\text{C10)} \quad \sqrt{j} \widehat{\text{Rem}}_j \equiv \sqrt{j} \text{Rem}^d_j(\hat{P}_j) \text{ converges to zero in probability as } j \rightarrow \infty.$$

The validity of the upper bound requires the following additional condition:

$$\text{C11)} \quad \sqrt{j} \left[\Psi^d_j(P_0) - \Psi(P_0) \right] \text{ converges to zero in probability as } j \rightarrow \infty.$$

For simplicity, we will take $\ell_n = 0$ in this section.

We now discuss the conditions. Condition C1) is an immediate consequence of C7) and $D^d(P)(o)$ being uniformly bounded in $P \in \mathcal{M}$, $d \in \mathcal{D}$, $o \in \mathcal{O}$. This will be plausible in many situations, including the examples in this paper. A more general Lindeberg-type condition also suffices [see Condition C1 in Luedtke and van der Laan, 2016], though we omit its presentation here for brevity.

The other three conditions all rely on terms like $\frac{1}{n} \sum_{j=0}^{n-1} R_j$ converging to zero in probability, possibly at some rate. Ideally we want a stochastic version of the fact that, for $\beta \in [0, 1)$,

$$\frac{1}{n} \sum_{j=1}^n j^{-\beta} \approx \frac{1}{n} \int_1^n j^{-\beta} dj \approx \frac{n^{-\beta}}{1-\beta} \text{ when } n \text{ is large.} \quad (\text{A.1})$$

Lemma 6 of Luedtke and van der Laan [2016] establishes this result. We restate it here for convenience.

Lemma A.1

(Lemma 6 in Luedtke and van der Laan, 2016). Suppose that R_j is some sequence of (finite) real-valued random variables such that $R_j = o_{P_0}(j^{-\beta})$ for some $\beta \in [0, 1)$, where we assume that each R_j is a function of $\{O_i: 1 \leq i \leq j\}$. Then,

$$\frac{1}{n} \sum_{j=0}^{n-1} R_j = o_{P_0}(n^{-\beta}).$$

Conditions C2) through C4) are now easily handled. Condition C2) is a consequence of the fact that

$$\frac{1}{n} \sum_{j=0}^{n-1} \left| \frac{\sigma_j^2}{\hat{\sigma}_j^2} - 1 \right| \leq \gamma^{-1} \frac{1}{n} \sum_{j=0}^{n-1} |\hat{\sigma}_j^2 - \sigma_j^2| \rightarrow 0 \text{ in probability as } n \rightarrow \infty,$$

where the inequality holds by C7) and the convergence holds by C9) Lemma A.1. Condition C9) is easily shown to hold under Glivenko-Cantelli conditions on the estimators $\hat{P}_j$ and d_j [see, e.g., Theorem 7 in Luedtke and van der Laan, 2016]. Conditions C3) and C4) are an immediate consequence of C10) and C11) combined with Lemma A.1.

While sufficient conditions for C11) should be developed in each individual example, we can give intuition as to why this condition should be reasonable. For any $P \in \mathcal{M}$, let $d(P)$ return a maximizer of (1). We are interested in ensuring that $\Psi^{d_n(P_0)} - \Psi^{d(P_0)}(P_0)$ is small, where d_n is our estimate of a maximizer of (1). This can be expected to hold when the parameter $P \mapsto \Psi^{d(P)}(P_0)$ has pathwise derivative zero at $P = P_0$, where the P_0 in the Ψ

argument is fixed. When well-defined, the pathwise derivative will be zero because $d(P)$ is chosen to maximize $\Psi^d(P_0)$ in d .

A.2 Efficiency when the maximizer in (1) is unique

We have presented a parametric-rate estimator for $\Psi_{\mathcal{H}}(P_0)$, but thus far we have not made any claims about the efficiency of our estimator. In this section, we consider a fixed parameter in (1) that does not rely on sample size. We therefore omit the n subscript in many quantities to indicate their lack of dependence on sample size. We will give conditions under which our estimator is asymptotically efficient among all regular, asymptotically linear estimators. The efficiency bound is not typically well-defined when the maximizer is non-unique due to the non-regularity of the problem - generally in this case no regular, asymptotically linear estimator exists, so neither does an efficient member of this class [Hirano and Porter, 2012]. Thus the conditions that we give in this section will typically only hold when the maximizer $d_0 \in \mathcal{D}$ in (1) is unique.

We use the following additional assumptions for our efficiency result:

- C5)** $E_{P_0} \left[\left(\hat{D}_j(O) - D^{d_0}(P_0)(O) \right)^2 \mid O_1, \dots, O_j \right] \rightarrow 0$ in probability as $j \rightarrow \infty$.
- C6)** There exists some $M < \infty$ such that $P_0 \left(D^{d_0}(P_0)(O) < M \right)$ and $P_0 \left(\hat{D}_j(O) < M \right)$ with probability approaching 1 as $j \rightarrow \infty$.
- C7)** $\inf_{j \geq 1} \hat{\sigma}_j^2 > \gamma$ with probability 1 over draws of $(O_j : j = 0, 1, \dots)$.

We discuss the conditions immediately following the theorem.

Theorem A.2—(Asymptotic efficiency). Suppose that Ψ does not depend on sample size and is pathwise differentiable with canonical gradient $D^{d_0}(P_0)$. Further suppose that $\ell_{\mathcal{H}} = o(n)$. If C1) through C7) hold, then

$$\bar{\sigma}_n^2 \rightarrow \text{Var}_{P_0} \left(D^{d_0}(P_0)(O) \right) \text{ in probability as } n \rightarrow \infty.$$

Furthermore,

$$\psi_n - \Psi(P_0) = \frac{1}{n} \sum_{i=1}^n D^{d_0}(P_0)(O_i) + o_{P_0}(n^{-1/2}).$$

Thus, ψ_n is asymptotically efficient among all regular, asymptotically linear estimators.

The proof is entirely analogous to the proof of Corollary 3 in Luedtke and van der Laan [2016] so is omitted. See Lemma 25.23 of van der Vaart [1998] for a proof of the fact that asymptotic linearity with the influence function given by the canonical gradient implies regularity.

The additional conditions needed for this result over Theorem 1 are mild when the maximizing index is unique. Condition C5) says that Ψ should have the same canonical gradient as Ψ^{d_0} . While this should be manually checked in each example, it will be fairly typical when the maximizer is unique, since in this case an arbitrarily small fluctuation of P_0 will generally not change the maximizer. This is similar to problems in introductory calculus where the derivative at the maximum is zero. Condition C5) requires that $\hat{D}^j(O)$ converge to $D^{d_0}(P_0)(O)$ in mean-squared error, which is to be expected if $\hat{P}_{nj}$ begins to approximate P_0 and d_{nj} converges to the unique maximizer d_0 as $n, j \rightarrow \infty$. Condition C6) is a bounding assumption on the canonical gradient and estimates thereof that will hold in many examples of interest. Finally, Condition C7) will hold if one knows that $\text{Var}_{P_0}[D^d(P)(O)]$ is bounded away from zero uniformly in $P \in \mathcal{M}$ and $d \in \mathcal{D}$, and uses this knowledge to truncate $\hat{\sigma}_j^2$ for some deterministic sequence $\gamma_j \rightarrow 0$. For γ_j sufficiently small and j sufficiently large this truncation scheme will then have no effect on the variance estimates $\hat{\sigma}_j^2$.

A.3 Computationally efficient implementation

There are several computationally efficient ways to compute our estimate. In Section 6.1 of Luedtke and van der Laan [2016], we show that the runtime of our estimator can be dramatically improved by running the algorithm used to compute each $\hat{P}_j$ a limited number of times, say ten times. We do not detail this approach here, though we note that the theorems we have presented are general enough to apply to this case.

An alternative approach to improve runtime is to use the estimator's online nature to compute it efficiently both in time and storage. Suppose that we have an algorithm to update the estimate $\hat{P}_{nj}$ of P_0 to the estimate $\hat{P}_{n(j+1)}$ based on the first j observations by looking at O_{j+1} only. This will often be feasible if the parameter of interest and the bias correction step only require estimates of certain components of P_0 , e.g. of a set of regression and classification functions. In these cases we can apply modern regression and classification approaches to estimate these quantities [see, e.g., Xu, 2011, Luts et al., 2014]. Often d_{nj} can also be obtained using online methods, and thus $\frac{1}{n-\ell_n} \sum_{j=\ell_n}^{n-1} \left[\Psi^{d_{nj}}(\hat{P}_{nj}) + \hat{D}_{nj}(O_{j+1}) \right]$ can be estimated online by keeping a running sum. This quantity is not equal to ψ_n because it does not yet include the weights.

It will not in general be possible to compute the weights online, though their computation does not require storing $O(n)$ observations in memory. We can estimate $\text{Var}_{P_0}(\hat{D}_{nj}(O))$

consistently using the r_j observations, where $r_j \rightarrow \infty$ but can grow very slowly (even $\log j$ suffices asymptotically, though such a slow growth is not recommended for finite samples). Given online estimates of these variances, it is then straightforward to compute both $\bar{\sigma}_n$ and the weights and incorporate these into our estimator. In some cases, we can compute the weights, and thus the estimate, in a truly online fashion. Describing general sufficient conditions for this appears to be difficult, but we conjecture that often this will not typically

hold if $\mathcal{D}_n$ is not of finite cardinality. The weights can be computed online in the maximal correlation example.

Appendix B McKeague and Qian [2015] example

B.1 Proofs and results

Lemma A.3

Fix $\tilde{\delta} > 0$ and $d \in \mathcal{D}_n$. For any P with $\min_k \in \mathcal{K}_n s_P^2(X_k) > \tilde{\delta}$ and $s_P^2(Y) > \tilde{\delta}$, (6) holds. Proof.

Straightforward but tedious calculations show that

$$\begin{aligned} \text{Rem}^d(P) = m & \left(\frac{1}{s_P(X_k)s_P(Y)} [s_P(X_k)s_P(Y) - s_0(X_k)s_0(Y)] [\text{Corr}_P(X_k, Y) - \text{Corr}_0(X_k, Y)] \right. \\ & \frac{(E_P[X_k] - E_{P_0}[X_k])(E_P[Y] - E_{P_0}[Y])}{s_P(X_k)s_P(Y)} \\ & - \frac{\text{Corr}_P(X_k, Y)}{2} \left[\frac{(E_P[X_k] - E_{P_0}[X_k])^2}{s_P^2(X_k)} + \frac{(E_P[Y] - E_{P_0}[Y])^2}{s_P^2(Y)} \right] \\ & \left. - \frac{\text{Corr}_P(X_k, Y)}{2s_P^2(X_k)s_P^2(Y)} [s_P(X_k)s_0(Y) - s_0(X_k)s_P(Y)]^2 \right). \end{aligned} \quad (\text{A.2})$$

The result follows by taking the absolute value of both sides, applying the triangle inequality, using that $ab \leq (a^2 + b^2)/2$ for any real a, b , $\text{Corr}_P(X_k, Y) \leq 1$, and the lower bound $\tilde{\delta}$ on the variances. $\square$

We now establish high probability bounds on the difference between $s_{P_j}^2(X_k)$, $s_{P_j}^2(Y)$, $\text{Corr}_{P_j}(X_k, Y)$ and $\hat{\sigma}_{nj}$ and their population counterparts, uniformly over $k \in \mathcal{K}_n$ and j . We will use $\lesssim$ to denote “less than or equal to up to a universal multiplicative constant”. Let $\mathcal{F}_n$ denote the following class of functions mapping from $\mathcal{O} \equiv \mathbb{R}^\infty \times \mathbb{R}$ to the real line:

$$\{(x, y) \mapsto x_k^r y^s : 0 \leq r, s \leq 4; r + s \leq 4; k \in \mathcal{K}_n\}. \quad (\text{A.3})$$

Note that $|\mathcal{F}_n| \lesssim p$. We will use this class to develop concentration results about our estimates the needed portions of the likelihood. This class is actually somewhat larger than is needed for most of our results, as in fact

$$\begin{aligned} & \left\{ (x, y) \mapsto x_k y : k \in \mathcal{K}_n \right\} \cup \left\{ (x, y) \mapsto x_k : k \in \mathcal{K}_n \right\} \cup \left\{ (x, y) \mapsto y \right\} \\ & \cup \left\{ (x, y) \mapsto x_k^2 : k \in \mathcal{K}_n \right\} \cup \left\{ (x, y) \mapsto y^2 \right\} \end{aligned}$$

suffices for concentrating our estimates of $\text{Corr}_0(X_k, Y)$, $s_0(X_k)$, and $s_0(Y)$. Nonetheless, using this larger class $\mathcal{F}_n$ will allow us to prove results about the concentration of $\hat{\sigma}_{nj}^2$ about σ_{nj}^2 and just stating it as a single class is convenient for brevity.

For $f \in \mathcal{F}_n$ and $j \in \{1, \dots, n\}$, define the empirical process as

$$\mathbb{G}_{nj} \equiv \frac{1}{\sqrt{j}} \sum_{i=1}^j [f(O_i) - P_0 f] = \sqrt{j}(P_j - P_0)f,$$

where we use P_j denote the empirical distribution of $O_0, \dots, O_{j-1}$ and $Pf \equiv E_P[f(O)]$ for any distribution P . Let $\|\mathbb{G}_{nj}\|_{\mathcal{F}_n} \equiv \sup_{f \in \mathcal{F}_n} |\mathbb{G}_{nj}|$. By Theorem 2.14.1 in van der Vaart and

Wellner [1996] shows that

$$E\|\mathbb{G}_{nj}\|_{\mathcal{F}_n} \lesssim \sqrt{\log \# \mathcal{F}_n} \lesssim \sqrt{\log p}, \quad (\text{A.4})$$

where the expectation is over the draws $O_1, \dots, O_j$. We have used that our class is bounded by the constant 1.

Let

$$K_{nj} \equiv j^{-1/2} \sqrt{\log \max\{n, p\}}. \quad (\text{A.5})$$

Define the events

$$\begin{aligned} \mathcal{A}_{nj} & \equiv \left\{ \max_{f \in \mathcal{F}_n} |(P_j - P_0)f| \leq CK_{nj} \right\} \text{ for all } j = 1, \dots, n, \\ \mathcal{A}_n & \equiv \bigcap_{j=1}^n \mathcal{A}_{nj}, \end{aligned}$$

where C in the definition of $\mathcal{A}_{nj}$ is equal the smallest universal constant satisfying (A.4) plus 1.

Lemma A.4

For any sample size n , the event $\mathcal{A}_n$ occurs with probability at least $1 - n/\max\{n^2, p\} - 1/n$.

Proof—We first upper bound the probability of the complement of $\mathcal{A}_{nj}$ for each n, j . Fix n and $j \leq n$. By the bounds on X and Y , changing one O_i in $(O_1, \dots, O_j)$ to some other value in the support of P_0 can change b by at most $1/\sqrt{j}$. Thus $(O_1, \dots, O_j) \mapsto \|\mathbb{G}_{nj}\|_{\mathcal{F}_n}$ satisfies the

bounded differences property with bound $1/\sqrt{j}$, and we may apply McDiarmid's inequality [McDiarmid, 1989] to show that, with probability at most $1 - \exp(-2t^2)$,

$$\|\mathbb{G}_{nj}\|_{\mathcal{F}_n} \leq E\|\mathbb{G}_{nj}\|_{\mathcal{F}_n} + t. \text{ Choosing } t = \sqrt{\frac{\log \max\{n^2, p\}}{2}} \text{ and using (A.4) yields that, with}$$

probability at least $1 - 1/\max\{n^2, p\}$, the following inequality holds for all $j = 1, \dots, n$:

$$\begin{aligned} \|\mathbb{G}_{nj}\|_{\mathcal{F}_n} &\leq E\|\mathbb{G}_{nj}\|_{\mathcal{F}_n} + \sqrt{\frac{\log \max\{n^2, p\}}{2}} \leq C'\sqrt{\log p} + \sqrt{\log \max\{n, p\}} \\ &\leq C\sqrt{\log \max\{n, p\}}, \end{aligned}$$

where C' denotes the universal constant in (A.4).

By DeMorgan's laws and a union bound, it follows that the event $\mathcal{A}_n \equiv \cap_j \mathcal{A}_{nj}$ occurs with probability at least $1 - n/\max\{n^2, p\} = 1 - 1/n$. $\square$

We have shown that $\mathcal{A}_n$ occurs with high probability. Now we show that our estimates of variances, covariances, and correlations perform well when $\mathcal{A}_n$ occurs.

Lemma A.5

Fix a sample size $n \geq 2$. The occurrence of $\mathcal{A}_n$ implies that, for all $j = 2, \dots, n$:

- 1) $\max_{k \in \mathcal{K}_n} |s_{P_j}(X_k) - s_0(X_k)| \lesssim \delta^{-1/2} K_{nj};$
- 2) $\max_{k \in \mathcal{K}_n} |s_{P_j}^2(X_k) - s_0^2(X_k)| \lesssim K_{nj};$
- 3) $|s_{P_j}(Y) - s_0(Y)| \lesssim \delta^{-1/2} K_{nj};$
- 4) $|s_{P_j}^2(Y) - s_0^2(Y)| \lesssim K_{nj};$
- 5) $\max_{k \in \mathcal{K}_n} |\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y)| \lesssim \delta^{-1} K_{nj},$

where we define $\text{Corr}_{P_j}(X_k, Y) = 0$ when either $s_{P_j}(X_k)$ or $s_{P_j}(Y)$ is equal to zero.

Proof—Suppose $\mathcal{A}_n$ holds and fix $k \in \mathcal{K}_n$. The triangle inequality and the bounds on X_k yield that

$$\begin{aligned}
|s_{P_j}^2(X_k) - s_0^2(X_k)| &= \left| E_{P_j}[X_k^2] - E_{P_0}[X_k^2] \right| - \left(E_{P_j}[X_k] + E_{P_0}[X_k] \right) \left(E_{P_j}[X_k] - E_{P_0}[X_k] \right) \\
&\leq |E_{P_j}[X_k^2] - E_{P_0}[X_k^2]| + 2|E_{P_j}[X_k] - E_{P_0}[X_k]| \lesssim K_{nj}.
\end{aligned}$$

This gives 2). For 1), note that

$$|s_{P_j}(X_k) - s_0(X_k)| = \left| \frac{s_{P_j}^2(X_k) - s_0^2(X_k)}{s_{P_j}(X_k) + s_0(X_k)} \right| \lesssim \delta^{-1/2} K_{nj}.$$

The same argument yields 3) and 4).

Again fix k . An application of the triangle inequality and the bounds on X_k and Y readily yield that $|\text{Cov}_{P_j}(X_k, Y) - \text{Cov}_{P_0}(X_k, Y)| \lesssim K_{nj}$. Furthermore,

$$\begin{aligned}
&\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y) \\
&= \frac{\text{Cov}_{P_j}(X_k, Y) - \text{Cov}_{P_0}(X_k, Y)}{s_0(X_k)s_0(Y)} - \frac{\text{Corr}_{P_j}(X_k, Y)s_{P_j}(Y)}{s_0(X_k)s_0(Y)} \left[s_{P_j}(X_k) - s_0(X_k) \right] \\
&\quad - \frac{\text{Corr}_{P_j}(X_k, Y)}{s_0(Y)} \left[s_{P_j}(Y) - s_0(Y) \right].
\end{aligned}$$

Taking the absolute value of both sides, applying the triangle inequality, and using the lower bounds on $s_0(X_k)$ and $s_0(Y)$ and the upper bound on $\text{Corr}_{P_j}(X_k, Y)$ yields that

$|\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y)| \lesssim \delta^{-1} K_{nj}$. This holds for all k , so 5) holds. $\square$

Lemma A.6

Let C be the smallest universal constant in 2) of that Lemma A.5, and let n be any natural number satisfying $n \geq \lceil 4C^{-2}\delta^{-2} \log \max\{n, p\} \rceil \equiv \mathcal{J}(n, \delta)$. Under these conditions, the occurrence of $\mathcal{A}_n$ implies that, for all $j = \mathcal{J}(n, \delta), \dots, n$,

$$8 \quad \min_{k \in \mathcal{K}_n} s_{P_j}^2(X_k) \geq \delta/2 \text{ and } \min_{k \in \mathcal{K}_n} s_{P_j}^2(Y) \geq \delta/2;$$

$$9 \quad \min_{k \in \mathcal{K}_n} \frac{s_{P_0}^2(X_k)}{s_{P_j}^2(X_k)} \leq \delta \text{ and } \frac{s_{P_0}^2(Y)}{s_{P_j}^2(Y)} \leq 2;$$

$$10 \quad |\hat{\sigma}_{nj}^2 - \sigma_{nj}^2| \lesssim \delta^{-2} K_{nj};$$

$$11 \quad |\hat{\sigma}_{nj}^2 - \text{Var}_{P_0}(D_{nj}(P_0)(O))| \lesssim \delta^{-2} K_{nj} + (\widehat{\text{Rem}}_{nj})^2.$$

Proof—By Lemma A.5, 2) holds, and using that $j \in \mathcal{J}(n, \delta)$, we see that

$$s_{P_j}^2(X_k) = s_{P_0}^2(X_k) + s_{P_j}^2(X_k) - s_{P_0}^2(X_k) \geq s_{P_0}^2(X_k) - \max_{k \in \mathcal{K}_n} |s_{P_j}^2(X_k) - s_{P_0}^2(X_k)| \geq \delta/2.$$

The same argument works for $s_{P_j}^2(Y)$, so 8 holds. Furthermore,

$$\frac{s_{P_0}^2(X_k)}{s_{P_j}^2(X_k)} \leq \frac{s_{P_0}^2(X_k)}{s_{P_0}^2(X_k) - CK_{nj}} = 1 + \frac{CK_{nj}}{s_{P_0}^2(X_k) - CK_{nj}} \leq 1 + 2C\delta^{-1}K_{nj} \leq 2,$$

where the final two inequalities hold by 8. This proves the first part of 9, and the bound on $s_{P_0}^2(Y)/s_{P_j}^2(Y)$ holds by the same argument. For the second result, note that

$$\begin{aligned} |\hat{\sigma}_{nj}^2 - \sigma_{jn}^2| &\leq |(P_j - P_0)\hat{D}_{nj}^2| + |(P_j\hat{D}_{nj})^2 - (P_0\hat{D}_{nj})^2| \\ &\leq |(P_j - P_0)\hat{D}_{nj}^2| + |(P_j + P_0)\hat{D}_{nj}| |(P_j - P_0)\hat{D}_{nj}|. \end{aligned}$$

Using 8, the bounds on X and Y , and the triangle inequality shows that

$$\lesssim |(P_j - P_0)\hat{D}_{nj}^2| + \delta^{-1} |(P_j - P_0)\hat{D}_{nj}| \lesssim \delta^{-2} \|\mathbb{G}_{nj}\|_{\mathcal{F}_n},$$

where we have used that $\mathcal{F}_n$ contains all polynomials of X_k, Y of degree at most 4. By the occurrence of $\mathcal{A}_n$, the final line is upper bounded by a constant times $\delta^{-2} K_{nj}$. This yields 10.

For 11, we will bound $|s_{nj}^2 - \text{Var}_{P_0}(D_{nj}(P_0)(O))|$ and then combine this with 10 using the triangle inequality. We have that

$$|s_{nj}^2 - \text{Var}_{P_0}(D_{nj}(P_0)(O))| \leq |P_0[\hat{D}_{nj}^2 - D_{nj}(P_0)^2]| + (P_0\hat{D}_{nj})^2.$$

Now we use that $P_0\hat{D}_{nj} = -\text{Corr}_{P_j}(X_k, Y) + \text{Corr}_{P_0}(X_k, Y) + \widehat{\text{Rem}}_{nj}$; and $(a + b)^2 \leq 2(a^2 + b^2)$

for any real a, b to see that $(P_0\hat{D}_{nj})^2 \lesssim \max_k \left(\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y) \right)^2 + (\widehat{\text{Rem}}_{nj})^2$. By 5)

from Lemma A.5 and the fact that $j \in \mathcal{J}(n, \delta)$, the maximum over $k \in \mathcal{K}_n$ is bounded above by a constant times $\delta^{-2} K_{nj}^2 \lesssim \delta^{-2} K_{nj}$. Continuing with the above,

$$\begin{aligned}
&\lesssim |P_0| \left(\left| \widehat{D}_{nj} + D_{nj}(P_0) \right| \left| \widehat{D}_{nj} - D_{nj}(P_0) \right| \right) + \delta^{-2} K_{nj} + \left(\widehat{\text{Rem}}_{nj} \right)^2 \\
&\lesssim \delta^{-1} P_0 | \widehat{D}_{nj} - D_{nj}(P_0) | + \delta^{-2} K_{nj} + \left(\widehat{\text{Rem}}_{nj} \right)^2 \\
&\lesssim \delta^{-2} \|G_{nj}\|_{\mathcal{F}_n} + \delta^{-2} K_{nj} + \left(\widehat{\text{Rem}}_{nj} \right)^2 \\
&\lesssim \delta^{-2} K_{nj} + \left(\widehat{\text{Rem}}_{nj} \right)^2,
\end{aligned}$$

where we used 8 for the second to last inequality. $\square$

Lemma A.7

Suppose the conditions of Lemma A.6. Under these conditions, the occurrence of $\mathcal{A}_n$ implies that, for all $j = \mathcal{J}(n, \delta), \dots, n$.

$$8. |\widehat{\text{Rem}}_{nj}| \lesssim \delta^{-5/2} \left(K_{nj} \right)^2.$$

Proof—By Lemma A.6, $\min_{k \in \mathcal{K}_n} s_{P_j}^2(X_k) \geq \delta/2$ and $\min_{k \in \mathcal{K}_n} s_{P_j}^2(Y) \geq \delta/2$. By Lemma A.3, this yields

$$\begin{aligned}
|\widehat{\text{Rem}}_{nj}| &\lesssim \delta^{-1} \max_{k \in \mathcal{K}_n} \left(|s_{P_j}(X_k)s_{P_j}(Y) - s_0(X_k)s_0(Y)| |\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_0(X_k, Y)| \right. \\
&\quad + \left(E_{P_j}[X_k] - E_{P_0}[X_k] \right)^2 + \left(E_{P_j}[Y] - E_{P_0}[Y] \right)^2 \\
&\quad \left. + \frac{s_0^2(Y)}{s_{P_j}^2(Y)} \left[s_{P_j}(X_k) - s_0(X_k) \right]^2 + \frac{s_0^2(X_k)}{s_{P_j}^2(X_k)} \left[s_{P_j}(Y) - s_0(Y) \right]^2 \right).
\end{aligned}$$

By the bounds on X and Y and the triangle inequality,

$$|s_{P_j}(X_k)s_{P_j}(Y) - s_0(X_k)s_0(Y)| \leq |s_{P_j}(X_k) - s_{P_0}(X_k)| + |s_{P_j}(Y) - s_{P_0}(Y)|.$$

Applying 9 from Lemma A.6 and the results of Lemma A.5 to the above yields the result. $\square$

Lemma A.8

Let γ be as defined in (7). For a constant $C(\gamma, \delta) > 0$ relying on γ and δ only, the occurrence of $\mathcal{A}_n$ implies that, for all $j = \lceil C(\gamma, \delta) \log \max\{n, p\} \rceil, \dots, n$.

$$8. \widehat{\sigma}_{nj}^2 \geq \gamma/2.$$

Sketch of proof—Suppose $\mathcal{A}_n$. By 11 and 8, for all $j \in \mathcal{J}(n, \delta)$

$$|\hat{\sigma}_{nj}^2 - \text{Var}_{P_0}(D_{nj}(P_0)(O))| \lesssim \delta^{-2} K_{nj} + \delta^{-10} (K_{nj})^4.$$

It is easy to confirm that, for a universal constant $C > 0$, the above yields that the left-hand side is upper bounded by $\gamma/2$ for all $j \leq C\gamma^{-1/2}\delta^{-2} \max\{\delta^{-3}, \gamma^{-3/2}\} \log \max\{n, p\} \equiv C(\gamma, \delta) \log \max\{n, p\} = \mathcal{J}(n, \delta)$. An application of the triangle inequality gives the result. $\square$

The remainder of the results in this section are asymptotic in nature. We omit the dependence on δ and γ in these statements as these quantities are treated as fixed as sample size grows. Throughout we assume that

$$\frac{\log \max\{n, p\}}{\ell_n} \rightarrow 0, \quad (\text{A.6})$$

$$\beta_n^2 \log \frac{n}{\ell_n} \rightarrow 0, \quad (\text{A.7})$$

$$\limsup_{n \rightarrow \infty} \frac{\ell_n}{n} < 1. \quad (\text{A.8})$$

In view of (A.6) and (A.7), we see that, roughly, ℓ_n grows faster than $\log \max\{n, p\}$ if β_n goes to zero faster than $1/\sqrt{\log n}$ and at least as fast as $n \exp(-o(\beta_n^{-2}))$ if β_n goes to zero more slowly than $1/\sqrt{\log n}$. Given an $\varepsilon > 0$, one possible choice of ℓ_n that satisfies these properties is

$$\ell_n = \max\left\{(\log \max\{n, p\})^{1+\varepsilon}, n \exp(-\beta_n^{-2+\varepsilon})\right\}.$$

We have the following result.

Lemma A.9

For all n large enough, ℓ_n is at least $\mathcal{J}(n, \delta)$ and is at least $C(\gamma, \delta) \log \max\{n, p\}$, where these quantities are defined in Lemmas A.6 and A.8, respectively.

Proof—This is an immediate consequence (A.6) of the fact that δ and γ are fixed as sample size grows. $\square$

Theorem A.10

C1), C2), and C3) hold.

Proof

- C1)** By Lemma A.9, we can apply 8 from Lemma A.6 and Lemma A.8 provided n is large enough. In that case, $\frac{\hat{D}_{nj}}{\hat{\sigma}_{nj}} \lesssim \delta^{-1} \gamma^{-1/2}$ for all $j \in \ell_n$ provided $\mathcal{A}_n$ holds. By Lemma A.4, this then occurs with probability at least $1 - 1/n$, and thus C1) holds.
- C2)** If $\mathcal{A}_n$ holds, then Lemmas A.8 and A.9 show that, for all n large enough,

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \left| \frac{\sigma_{nj}^2}{\hat{\sigma}_{nj}^2} - 1 \right| \leq \frac{2\gamma^{-1}}{n - \ell_n} \sum_{j=\ell_n}^{n-1} |\hat{\sigma}_{nj}^2 - \sigma_{nj}^2| \lesssim \frac{\gamma^{-1} \delta^{-2}}{n - \ell_n} \log \max \left\{ n, p \right\} \sum_{j=\ell_n}^{n-1} j^{-1}$$

By 10 in Lemma A.6 and the fact that $\sum_{a-1}^b j^{-1} \leq \int_{a-1}^b j^{-1} dj$, the right-hand side is has an upper bound proportional to $\frac{\gamma^{-1} \delta^{-2}}{n - \ell_n} \log \max \left\{ n, p \right\} \log n$. This bound is $\alpha(1)$ by (A.8) and the fact that $\beta_n \rightarrow 0$. The fact that $\mathcal{A}_n$ occurs with probability approaching 1 (Lemma A.4) yields C2).

- C3)** Suppose that n is large enough so that the results of Lemma A.9 apply. Also suppose that $\mathcal{A}_n$ occurs. We have that

$$\begin{aligned} \frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \frac{\widehat{\text{Rem}}_{nj}}{\hat{\sigma}_{nj}} &\lesssim \frac{\gamma^{-1/2} \delta^{-5/2}}{n - \ell_n} \sum_{j=\ell_n}^{n-1} (K_{nj})^{20} && \text{(Lemmas A.7, A.8, and A.9)} \\ &= \frac{\gamma^{-1/2} \delta^{-5/2}}{n - \ell_n} \log \max \left\{ n, p \right\} \sum_{j=\ell_n}^{n-1} j^{-1} && \text{(Eq. A.5)} \\ &\lesssim \frac{\gamma^{-1/2} \delta^{-5/2}}{n - \ell_n} \log \max \left\{ n, p \right\} \log \frac{n}{\ell_n} && \left(\sum_{a-1}^b j^{-1} \leq \int_{a-1}^b j^{-1} dj \right) \\ &= o \left([n - \ell_n]^{-1/2} \right). && \text{(Eqs. A.7 and A.8)} \end{aligned}$$

The fact that $\mathcal{A}_n$ occurs with probability approaching 1 (Lemma A.4) yields C3). $\square$

Let k_{n0} be a possibly non-unique k maximizer of $|\text{Corr}_{P_0}(X_k, Y)|$. For each $r > 0$, let $\mathcal{K}_n^r \subseteq \mathcal{K}_n$ denote the set of all $k \in \mathcal{K}_n$ such that $|\text{Corr}_{P_0}(X_{k_{n0}}) - \text{Corr}_{P_0}(X_k, Y)| \leq r$.

The upcoming theorem uses the following conditions to establish the validity of a hypothesis test of no effect and of the upper bound of our confidence interval, respectively:

M1) For some sequence $\{t_n\}$ with $t_n \rightarrow +\infty$, there exists a sequence of non-empty subsets $\mathcal{K}_n^r \subseteq \mathcal{K}_n$ such that, for all n ,

$$\inf_{k_1 \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_{k_1}, Y)| \geq \sup_{k_2 \in \mathcal{K}_n \setminus \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_{k_2}, Y)| + t_n n^{-1/4} \beta_n.$$

If $\mathcal{K}_n^r = \mathcal{K}_n$, then the supremum on the right-hand side is taken to be zero.

M2) The conditions of M1) hold, and also

$$\text{Diam}(\mathcal{K}_n^*) \equiv \sup_{k_1, k_2 \in \mathcal{K}_n^r} \left(|\text{Corr}_{P_0}(X_{k_1}, Y)| - |\text{Corr}_{P_0}(X_{k_2}, Y)| \right) = o(n^{-1/2}).$$

The first of these conditions will be used to establish the consistency of a null hypothesis significance test. The second of these conditions is similar to margin conditions used in classification, and will be used to establish the validity of our confidence interval.

Theorem A.11

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \hat{\sigma}_{nj}^{-1} \left[\Psi^{dj}(P_0) - \Psi_n(P_0) \right] = O_{P_0}(n^{-1/4} \beta_n). \quad (\text{A.9})$$

If also M1), then the right-hand side of the above can be tightened to $O_{P_0}(\text{Diam}(\mathcal{K}_n^*) \wedge n^{-1/4} \beta_n) + O_{P_0}(n^{1/2})$. If also M2), then C4) holds.

Proof—Suppose that $\mathcal{A}_n$ holds and n is large enough so that the results of Lemma A.9 apply. For each $j \in \ell_n$ let k_{nj} represent the $k \in \mathcal{K}_n$ which maximizes $|\text{Corr}_{P_j}(X_k, Y)|$. Let $m_0 = \text{sgn}[\text{Corr}_{P_0}(X_{k_{n0}}, Y)]$ and $m_{nj} = \text{sgn}[\text{Corr}_{P_j}(X_{k_{nj}}, Y)]$. Then, for a universal constant $C > 0$,

$$\begin{aligned}
0 &\geq m_0 \text{Corr}_{P_j}(X_{k_{n0}}, Y) - |\text{Corr}_{P_j}(X_{k_{nj}}, Y)| \\
&= m_0 \left[\text{Corr}_{P_0}(X_{k_{n0}}, Y) - m_{nj} \text{Corr}_{P_0}(X_{k_{nj}}, Y) \right] \\
&\quad + m_0 \left[\text{Corr}_{P_j}(X_{k_{n0}}, Y) - \text{Corr}_{P_0}(X_{k_{n0}}, Y) \right] - m_{nj} \left[\text{Corr}_{P_j}(X_{k_{nj}}, Y) - \text{Corr}_{P_0}(X_{k_{nj}}, Y) \right] \\
&\geq \Psi_n(P_0) - \Psi^{d_{nj}}(P_0) - 2 \max_{k \in \mathcal{K}_n} |\text{Corr}_{P_j}(X_k, Y) - \text{Corr}_{P_0}(X_k, Y)| \\
&\geq \Psi_n(P_0) - \Psi^{d_{nj}}(P_0) - C\delta^{-1}K_{nj},
\end{aligned} \tag{A.10}$$

where the final inequality holds by Lemma 5). Using that $\sum_{j=\ell_n}^{n-1} j^{-1/2} \lesssim \sqrt{n}$ and (A.8),

$$\frac{1}{n-\ell_n} \sum_{j=\ell_n}^{n-1} K_{nj} \lesssim \log \max\{n, p\} \frac{\sqrt{n}}{n-\ell_n} \lesssim n^{-1/4} \beta_n.$$

By Lemma A.8, this then implies that the left-hand side of (A.9) is upper bounded by an $O(\gamma^{-1/2} n^{-1/4} \beta_n)$ term under $\mathcal{A}_n$, and so Lemma A.4 yields (A.9).

For the second result, suppose that M1) holds. Observe that, for all $j > Cnt_n^{-1}$ for C as

defined in (A.10), $\Psi_n(P_0) - \Psi^{d_{nj}}(P_0) < t_n n^{-1/4} \beta_n$. Furthermore,

$|\Psi_n(P_0) - |\text{Corr}_{P_0}(X_{k_{nj}})| \leq \Psi_n(P_0) - \Psi^{d_{nj}}(P_0)$. Thus $k_{nj} \in \mathcal{K}_n^*$ as defined in M1). Furthermore, m_{nj} must equal $\text{sgn}[\text{Corr}_{P_0}(X_{k_{nj}}, Y)]$, since otherwise

$$\Psi_n(P_0) - \Psi^{d_{nj}}(P_0) \geq |\text{Corr}_{P_0}(X_{k_{nj}}, Y)| - \Psi^{d_{nj}}(P_0) \geq 2 \inf_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_{k_{nj}}, Y)| \geq 2t_n n^{-1/4} \beta_n,$$

contradicting the fact that $\Psi_n(P_0) - \Psi^{d_{nj}}(P_0) \leq t_n n^{-1/4} \beta_n$ per (A.10). Because $k \in \mathcal{K}_n^*$, we see

that $\Psi^{d_{nj}}(P_0) \geq \inf_{k \in \mathcal{K}_n^*} |\text{Corr}_{P_0}(X_{k_{nj}}, Y)|$. Hence,

$$\frac{1}{n-\ell_n} \sum_{j=\max\{\ell_n, \lceil Cnt_n^{-1} \rceil\}}^{n-1} \left[\Psi^{d_{nj}}(P_0) - \Psi_n(P_0) \right] \geq -\text{Diam}(\mathcal{K}_n^*). \tag{A.11}$$

Further, if $\lceil Cnt_n^{-1} \rceil \geq \ell_n$, (A.10) yields

$$\left[\sum_{j=\ell_n}^{Cn\ell_n^{-1}} \left[\Psi^{dj}(P_0) - \Psi_n(P_0) \right] \right] \geq -C \sum_{j=\ell_n}^{Cn\ell_n^{-1}} K_{nj} \geq -C \log \max\{n, p\} \int_{\ell_n^{-1}}^{Cn\ell_n^{-1}} j^{-1/2} dj$$

It follows that the left-hand side above is greater than or equal to a positive universal constant times $n^{1/2} t_n^{-1/2}$. Dividing the left by $n - \ell_n$ and applying (A.8) yields that this same result holds with an upper bound on the order of $n^{-1/2} t_n^{-1/2}$. Combining this with (A.11) shows that

$$\frac{1}{n - \ell_n} \sum_{j=\ell_n}^{n-1} \left[\Psi^{dj}(P_0) - \Psi_n(P_0) \right] \geq -\text{Diam}(\mathcal{K}_n^*) + O(n^{-1/2} t_n^{-1/2}).$$

Using that $t_n^{-1/2} \rightarrow 0$, $n^{-1/2} t_n^{-1/2} = o(n^{-1/2})$. When proving the first result (A.9) we also showed that the left-hand side is upper-bounded by a positive constant times $-\delta^{-1} n^{-1/4} \beta_n$. Combining with Lemma A.8 and using that $\mathcal{A}_n$ holds with probability approaching 1 (Lemma A.4) shows that the left-hand side of (A.9) is $O_P(\text{Diam}(\mathcal{K}_n^*) \wedge n^{-1/4} \beta_n) + o_P(n^{-1/2})$. If M2) holds, then this expression is $o_P(n^{-1/2})$, and so C4) holds. $\square$

B.2 Computationally efficient implementation of our estimator

In this section, we describe how to implement the estimator for the McKeague and Qian [2015] example in $O(np)$ time. We show that this can be accomplished using $O(p)$ storage when the observations $O_1, \dots, O_n$ arrive in a stream.

Fix n so that the set $\mathcal{K}_n$ of predictor indices is also fixed. For each j , let P_j denote the empirical distribution of the first j observations. Recall the definition of the class $\mathcal{F}_n$ from (A.3), and note that $\mathcal{F}_n$ contains $O(p)$ functions. It is easy to see that, at $j=2$, we can compute $P_j f \equiv E_{P_j}[f(O)]$ for each $f \in \mathcal{F}_n$ using $O(p)$ time and storage. Furthermore, for $j \geq 3$ the fact that $P_j f = f(O_j) + \frac{j-1}{j} P_{j-1} f$ shows that we can compute and save $P_j f$ in $O(p)$ time and storage if we know O_j and $P_{j-1} f$. To attain this storage complexity, we remove $P_{j-2} f$, $f \in \mathcal{F}_n$, from memory for each $j \geq 4$ so that $P_2 f, \dots, P_{j-2} f$ are not stored in memory.

We now have an algorithm that, at observation j , starts with O_j and $P_{j-1} f$, $f \in \mathcal{F}_n$, stored in memory and, after running the steps described in the preceding paragraph, also has $P_j f$, $f \in \mathcal{F}_n$ stored in memory. Given $P_j f$, $f \in \mathcal{F}_n$, one can compute and save $\text{Cov}_{P_j}(X_k, Y) = E_{P_j}[X_k Y] - E_{P_j}[X_k] E_{P_j}[Y]$, $k \in \mathcal{K}_n$, and $s_{P_j}^2(Z) = E_{P_j}[Z^2] - E_{P_j}[Z]^2$, Z equal to Y or X_k , $k \in \mathcal{K}_n$, in $O(p)$ time and storage. We can now compute and save

$\text{Corr}_{P_j}(X_k, Y) = \frac{\text{Cov}_{P_j}(X_k, Y)}{s_{P_j}(X_k)s_{P_j}(Y)}$, $k \in \mathcal{K}_n$ in $O(p)$ time and storage. If the predictors or outcome

are large and their variance small, the described online computation of the sample variance may lead to numerical difficulties. See Welford [1962] for a better estimate of the variance in this setting.

Let H_j denote the collection of (i) the integer j , (ii) $P_j f$, $f \in \mathcal{F}_n$, (iii) $s_{P_j}^2$, and (iv) $\text{Cov}(X_k, Y)$, $s_{P_j}^2(X_k)$ and $\text{Corr}_{P_j}(X_k, Y)$, $k \in \mathcal{K}_n$. For $j = 2$, let UPDATEH be a function which takes as input (O_{j+1}, H_j) and outputs H_{j+1} . We have shown that UPDATEH(O_{j+1}, H_j) can run in $O(p)$ time for any $j = 2$. We call a separate function INITIALIZEH on (O_1, O_2) to obtain the initial value H_2 . This function runs in $O(p)$ time and storage.

Let the function MAXIMIZER be a function that takes as input H_j and returns the $d_j = (k_j, m_j)$ which maximizes $m\text{Corr}_{P_j}(X_k, Y)$ in $d \equiv (k, m) \in \mathcal{K}_n$, thereby allowing us to compute

$\hat{\sigma}_{nj} = P_j D^{d_j}(P_j)^2$. Finding d_j involves finding the maximum of $|\mathcal{D}_n| = 2p$ numbers, and therefore can be accomplished in $O(p)$ time.

The function CALCD takes as input H_j , O_{j+1} , and d_j and calculates $D^{d_j}(P_j)(O_{j+1})$. It is easy to see that this can be accomplished in $O(1)$ time and $O(p)$ storage.

For ease of notation in the proceeding paragraph and equation we omit the dependence of $d_j = (k_j, m_j)$ on j in the notation. Since $D^d(P_j)$ is a gradient for Ψ^d at P_j and gradients are mean zero, $P_j D^d(P_j) = 0$. For any $d \in \mathcal{D}_n$, tedious but trivial calculations show that

$$\begin{aligned} P_j D^d(P_j)^2 &= \left[\frac{2 + \text{Corr}_{P_j}(X_k, Y)^2}{2s_{P_j}^2(X_k)s_{P_j}^2(Y)} \sum_{r=0}^2 \sum_{s=0}^2 (-1)^{r+s} \binom{2}{r} \binom{2}{s} E_{P_j}[X_k^r Y^s] E_{P_j}[X_k]^{2-r} E_{P_j}[Y]^{2-s} \right. \\ &\quad + \frac{\text{Corr}_{P_j}(X_k, Y)^2}{4} \sum_{r=0}^4 (-1)^r \binom{4}{r} \left[\frac{E_{P_j}[X_k^r] E_{P_j}[X_k]^{4-r}}{s_{P_j}^4(X_k)} + \frac{E_{P_j}[Y^r] E_{P_j}[Y]^{4-r}}{s_{P_j}^4(Y)} \right] \\ &\quad - \frac{\text{Corr}_{P_j}(X_k, Y)}{s_{P_j}(X_k)^3 s_{P_j}(Y)} \sum_{r=0}^3 \sum_{s=0}^1 (-1)^{r+s} \binom{3}{r} \binom{1}{s} E_{P_j}[X_k^r Y^s] E_{P_j}[X_k]^{3-r} E_{P_j}[Y]^{1-s} \\ &\quad \left. - \frac{\text{Corr}_{P_j}(X_k, Y)}{s_{P_j}(X_k) s_{P_j}(Y)^3} \sum_{r=0}^1 \sum_{s=0}^3 (-1)^{r+s} \binom{1}{r} \binom{3}{s} E_{P_j}[X_k^r Y^s] E_{P_j}[X_k]^{1-r} E_{P_j}[Y]^{3-s} \right]. \end{aligned}$$

Observe that all expectations on the right-hand side above are expectations over some $f \in \mathcal{F}_n$ applied to the observed data structure. It follows that the above can be computed in $O(1)$ time using a subset of the $O(p)$ expectation, standard deviation, and correlation

estimates stored in H_j . Let CALCSIGHAT denote the function which takes as input H_j and d_j and outputs $\hat{\sigma}_{nj}$. We have shown that CALCSIGHAT (H_j, d_j) runs in $O(1)$ time.

The pseudocode in ESTPSI describes our estimator, with most of the work done in the recursion step described in the function RECURSION. Because each call of RECURSION runs in $O(p)$ time, the $n - \ell_n = O(n)$ step for loop in ESTPSI requires time $O(np)$ time. The storage requirement of each call of RECURSION is $O(p)$. Because the code in the for loop in ESTPSI deletes the output from the previous recursion step, the total storage requirement of ESTPSI is $O(p)$.

Algorithm Recursion Step for Estimating $\Psi(P_0)$

```

function Recursion( $O_{j+1}, \psi_j, H_j, \bar{\sigma}_j, \ell_n$ )

2:   if  $j < \ell_n$  then  $\psi_{j+1} = 0$  and  $\bar{\sigma}_{j+1} = 0$ 
      else
4:      $d_j = \text{Maximizer}(H_j)$ 
         $\hat{\sigma}_{nj} = \text{CalcSigHat}(H_j, d_j)$ 
6:      $D^{j(P_j)}(O_{j+1}) = \text{CalcD}(H_j, O_{j+1})$ 
         $\psi_{j+1} = \frac{\psi_j}{\bar{\sigma}_j} + \frac{\text{Corr}_{P_j}(X_{d_j}, Y) + D^{j(P_j)}(O_{j+1})}{\hat{\sigma}_{nj}}$ 
         $\bar{\sigma}_{j+1} = \frac{1}{j+1} [j\bar{\sigma}_j + \hat{\sigma}_{nj}]$ 
         $H_{j+1} = \text{UpdateH}(O_{j+1}, H_j)$ 
10:    return  $(\psi_{j+1}, \bar{\sigma}_{j+1}, H_{j+1})$ 

```

► By convention, $0/0 = 0$.

Algorithm Estimate $\Psi(P_0)$ Using Sample of Size n

```

function ESTPSI( $n, \ell_n$ )

2:   Read  $O_1, O_2$  from data stream

      Base case:  $\psi_2 = 0, \bar{\sigma}_2 = 0$ , and  $H_2 = \text{InitializeH}(O_1, O_2)$ 

      for  $j = 2, \dots, n-1$  do
4:     Read  $O_{j+1}$  from data stream

         $(\psi_{j+1}, \bar{\sigma}_{j+1}, H_{j+1}) = \text{Recursion}(O_{j+1}, \psi_j, H_j, \bar{\sigma}_j, \ell_n)$ 
6:     Remove  $(O_{j+1}, \psi_j, H_j, \bar{\sigma}_j)$  from memory

```

8: **return** Point estimate ψ_n and confidence interval $\left[\psi_n \pm 1.96\bar{\sigma}_n/\sqrt{n}\right]$

References

- Bickel PJ, Klaassen CAJ, Ritov Y, Wellner JA. Efficient and adaptive estimation for semiparametric models. Johns Hopkins University Press; Baltimore: 1993.
- Chakraborty B, Moodie EE. Statistical Methods for Dynamic Treatment Regimes. Springer; Berlin Heidelberg New York: 2013.
- Gaenssler P, Strobel J, Stute W. On central limit theorems for martingale triangular arrays. Acta Math Hungar. 1978; 31(3):205–216.
- Hirano K, Porter JR. Impossibility results for nondifferentiable functionals. Econometrica. 2012; 80(4):1769–1790.
- Luedtke AR, van der Laan MJ. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. Annals of Statistics. 2016; 44(2):713–742.
- Luts J, Broderick T, Wand MP. Real-time semiparametric regression. Journal of Computational and Graphical Statistics. 2014; 23(3):589–615.
- McDiarmid C. On the method of bounded differences. Surveys in combinatorics. 1989; 141(1):148–188.
- McKeague IW, Qian M. An adaptive resampling test for detecting the presence of significant predictors. Journal of the American Statistical Association. 2015; 110(512)
- Pfanzagl J. Estimation in semiparametric models. Springer; Berlin Heidelberg New York: 1990.
- R Core Team. R: a language and environment for statistical computing. 2014. URL <http://www.r-project.org/>
- van der Laan MJ, Lendle SD. Online Targeted Learning. Division of Biostatistics, University of California; Berkeley: 2014. Technical Report 330 available at <http://www.bepress.com/ucbbiostat>
- van der Laan MJ, Robins JM. Unified methods for censored longitudinal data and causality. Springer; New York Berlin Heidelberg: 2003.
- van der Vaart AW. On differentiable functionals. Annals of Statistics. 1991; 19:178–204.
- van der Vaart AW. Asymptotic statistics. Cambridge University Press; New York: 1998.
- van der Vaart AW, Wellner JA. Weak convergence and empirical processes. Springer; Berlin Heidelberg New York: 1996.
- Welford BP. Note on a method for calculating corrected sums of squares and products. Technometrics. 1962; 4(3):419–420.
- Xu W. Towards optimal one pass large scale learning with averaged stochastic gradient descent. arXiv preprint arXiv:1107.2490. 2011
- Zhang Y, Laber EB. Comment. J Am Stat Assoc. 2015; 110(512):1451–1454.

Figure 1.

Power of the various testing procedures for (n, p) equal to (200, 200) and (500, 2000) under homoscedastic errors. The ART and ZL procedure performs the best in this setting.

Figure 2.

Power of the various testing procedures for (n, p) equal to $(200, 200)$ and $(500, 2000)$ under heteroscedastic errors. The ZL procedure fails to control the type I error in this setting.

Figure 3.

Power of the test from the stabilized one-step and from the Bonferroni-adjusted t -test for (n, p) equal to (2 000, 30 000) under homoscedastic errors.