NBER WORKING PAPER SERIES

SEMIPARAMETRIC ESTIMATION OF
TREATMENT EFFECTS IN RANDOMIZED EXPERIMENTS

Susan Athey
Peter J. Bickel
Aiyou Chen
Guido Imbens
Michael Pollmann

Working Paper 29242
<http://www.nber.org/papers/w29242>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2021

Generous support from the Office of Naval Research through ONR grant N00014-17-1-2131 (Athey and Imbens) is gratefully acknowledged. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Susan Athey, Peter J. Bickel, Aiyou Chen, Guido Imbens, and Michael Pollmann. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Semiparametric Estimation of Treatment Effects in Randomized Experiments
Susan Athey, Peter J. Bickel, Aiyou Chen, Guido Imbens, and Michael Pollmann
NBER Working Paper No. 29242
September 2021
JEL No. C01,C1,C14

ABSTRACT

We develop new semiparametric methods for estimating treatment effects. We focus on a setting where the outcome distributions may be thick tailed, where treatment effects are small, where sample sizes are large and where assignment is completely random. This setting is of particular interest in recent experimentation in tech companies. We propose using parametric models for the treatment effects, as opposed to parametric models for the full outcome distributions. This leads to semiparametric models for the outcome distributions. We derive the semiparametric efficiency bound for this setting, and propose efficient estimators. In the case with a constant treatment effect one of the proposed estimators has an interesting interpretation as a weighted average of quantile treatment effects, with the weights proportional to (minus) the second derivative of the log of the density of the potential outcomes. Our analysis also results in an extension of Huber's model and trimmed mean to include asymmetry and a simplified condition on linear combinations of order statistics, which may be of independent interest.

Susan Athey
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
athey@stanford.edu

Guido Imbens
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
Imbens@stanford.edu

Peter J. Bickel
University of California, Berkeley
Department of Statistics
411 Evans Hall
Berkeley, CA 94720
bickel@stat.berkeley.edu

Michael Pollmann
Stanford University
pollmann@stanford.edu

Aiyou Chen
Google
Mountain View, CA
aiyouchen@google.com

1. Introduction

Historically, randomized experiments were often carried out in medical and agricultural settings. In these settings, sample sizes were often modest, typically on the order of hundreds or (occasionally) thousands of units. Outcomes commonly studied included average mortality or crop yield, and were characterized by relatively well-behaved data with thin tails. Standard analyses in those settings typically involved estimating the average effect of the treatment by the difference in average outcomes by treatment group, followed by constructing confidence intervals for the average causal effect using Normal distribution based approximations. These methods originated in the 1920s, e.g., Neyman (1990); Fisher (1937), but they continue to be the standard in modern applications. See Wu and Hamada (2011) for a recent discussion.

More recently many experiments are conducted using digital technology (see, e.g., Kohavi et al. (2020) for an overview). Gupta et al. (2019, p.20) claim that “Together these organizations [Airbnb, Amazon, Booking.com, Facebook, Google, LinkedIn, Lyft, Microsoft, Netflix, Twitter, Uber, and Yandex] tested more than one hundred thousand experiment treatments last year.” The settings for these online experiments are substantially different from those in biomedical and agricultural settings. First, the experiments are often on a vastly different scale, with the number of units on the order of hundreds of thousands to tens of millions. Second, the outcomes of interest, variables such as time spent by a consumer, sales per consumer or payments per service provider, are characterized by distributions with extremely thick tails. Third, the average treatment effects are often small relative to the standard deviation of the outcomes, even if their magnitude is substantively important.

For example, Lewis and Rao (2015) analyze challenges with statistical power in experiments designed to measure digital advertising effectiveness. They discuss a hypothetical experiment where the average expenditure per potential customer is \$7, with a standard deviation of \$75, and where an average treatment effect of \$0.35 (0.005 of a standard deviation) would be substantial in the sense of being highly profitable for the company. In that example, an experiment with power for a treatment effect of \$0.35 equal to 80%, and a significance level for the test of means of 0.05, would require a sample size of 1.4 million customers. As a result confidence intervals for the average effect of an advertisement are likely to include zero even if the average effect were substantively important, even with large sample sizes, *e.g.*, over a hundred thousand units. However, the fact that a confidence interval for the average treatment effect includes zero does not mean that the data are not informative about the presence of causal effects of the treatment. Using Fisher exact p-value calculations (Fisher, 1937) with well-chosen statistics (*e.g.*, the Hodges-Lehman difference in average ranks, Rosenbaum (1993)), one may well be able to establish conclusively that treatment effects are present. However, the magnitude of the treatment effect is important for decisions about advertising. In this type of

example, point estimates of average effects are of more importance for decision makers than rejections of the sharp Fisher null hypothesis of no effects whatsoever.

This sets the stage for the problem we address in this paper. In the absence of additional information, there is no estimator for the average treatment effect that is more efficient than the difference in means. To obtain more precise estimates, we either need to change focus away from the average treatment effect, or we need to make additional assumptions. One approach to changing the question slightly is to transform the outcome (*e.g.* taking logarithms or winsorizing the outcome) followed by a standard analysis estimating the average effect of the treatment on the transformed outcome. In this paper we choose a different approach, making additional assumptions on the distribution of the data. The key conceptual contribution is that we postulate low-dimensional parametric models for the unit-level treatment effects. Note that we do not use a parametric models for the full outcome distributions, which might involve interaction between potential treatment and control values because specifying a parametric model that well approximates the full outcome distribution is more challenging than doing so for treatment effects. Unlike outcomes, treatment effects tend to be small and often have little variation. Instead, we propose a semiparametric model of the outcome distribution that incorporates a parametric model only for the treatment effects but not for the joint outcome distributions (Bickel et al. (1993), Bickel and Doksum (2015)). For this semiparametric set-up, we derive the influence function, the semiparametric efficiency bound, and propose semiparametrically efficient estimators. It turns out that the parametrization of the treatment effect can be very informative, potentially making the asymptotic variance for these semiparametric estimators much smaller than the corresponding asymptotic variance for the difference-in-means estimator. In addition, even if the model for the treatment effect is misspecified, the estimand corresponding to the estimator may have an interesting, causal, interpretation.

The remainder of the paper is organized as follows. First, in Section 2 we consider a special case where we assume that the treatment effect is additive and constant. This implies that the two potential outcome distributions differ only by a shift. This case is closely related to the classical two-sample problem, as discussed in Hodges Jr and Lehmann (1963), and to problems considered in the literature on descriptive statistics for nonparametric models as in Bickel and Lehmann (1975a,b), Doksum (1974), Doksum and Sievers (1976). Particularly relevant is the work of Jaeckel (1971a), Jaeckel (1971b). One way to interpret Section 2 is that it presents, in a causal inference framework, extensions of Jaeckel’s work to the two sample context in the setting of semiparametric theory. In the process of doing so, we generalize Huber’s model (Huber, 1964) and the estimator based on trimmed means to include asymmetry, and present a simplified version of the results of Chernoff et al. (1967) (see also Bickel (1967), Govindarajulu et al. (1967) and Stigler (1974) on linear combinations of order statistics). In particular,

we exhibit efficient M and L estimates for outcome distributions that are known up to a shift. We then analyze partially adaptive estimates, in particular flexible trimmed means (as in Jaeckel (1971a)) and fully adaptive estimates of both types, as discussed, for example, in Bickel et al. (1993). In this setting the problem is closely related to the literature on robust estimation of locations (*e.g.*, Bickel and Lehmann (1975a,b); Jaeckel (1971a); Hampel et al. (2011); Huber (2011); Bickel and Lehmann (1976, 2012)). In Section 3 we consider the case where we have more flexible parametric models for the treatment effect. In Section 4 we provide some simulation evidence regarding the finite sample properties of the proposed methods in controlled settings. In this section we also provide two applications of the methods, one to data on housing prices and one to data on medical expenses. In both cases the outcome distributions are thick-tailed. Section 5 concludes. A software implementation for R is available at <https://github.com/michaelpollmann/parTreat>.

2. Constant Additive Treatment Effects

In this section we focus on a special case of the general problem where the treatment effect is additive and constant across units. In Section 3 we discuss the general case. We start by setting up the problem more formally and introducing some notation. Then we discuss robust estimation in the one-sample case to motivate a class of weighted quantile treatment effect estimators. We then discuss the formal semiparametric problem and show the adaptivity of the proposed estimators. Finally we consider partial adaptivity and robustness.

2.1. Set Up

We consider a set up with a randomized experiment with n observations drawn randomly from a large population. With probability $p \in (0, 1)$ a unit is assigned to the treatment group. Let n_1 and $n_0 = n - n_1$ denote the number of units assigned to the treatment and control group. Using the potential outcome set up (Neyman, 1990; Rubin, 1974; Imbens and Rubin, 2015), let $Y_i(0)$ and $Y_i(1)$ denote the two potential outcomes for unit i , and let the binary treatment be denoted by $Z_i \in \{0, 1\}$. The realized outcome is $Y_i = Y_i(Z_i)$. We assume that there is no interference between units, so that the treatment assignment for one unit does not affect the outcomes for any other unit. For all units in the sample we observe the pair (Z_i, Y_i) .

The cumulative distribution functions for the two potential outcomes are $F_0(\cdot)$ and $F_1(\cdot)$ with inverses $F_0^{-1}(s)$ and $F_1^{-1}(s)$ and means and variances μ_0, μ_1, σ_0^2 and σ_1^2 . Note that by the random assignment assumption the distribution of the potential outcome $Y_i(z)$ is identical to the conditional distribution of the realized outcome Y_i conditional on $Z_i = z$.

[4] *Athey et al.*

We are interested in the average treatment effect in the population,

$$\tau = \mathbb{E}[Y_i(1) - Y_i(0)]. \quad (2.1)$$

The natural estimator for this average treatment effect is the difference in sample averages

$$\hat{\tau} = \bar{Y}_1 - \bar{Y}_0, \quad \text{where } \bar{Y}_1 = \frac{1}{n_1} \sum_{i=1}^n Z_i Y_i, \quad \bar{Y}_0 = \frac{1}{n_0} \sum_{i=1}^n (1 - Z_i) Y_i, \quad (2.2)$$

are the averages of the observed outcomes by treatment group. Under standard conditions

$$\sqrt{n}(\hat{\tau} - \tau) \xrightarrow{d} \mathcal{N}\left(0, \frac{\sigma_0^2}{1-p} + \frac{\sigma_1^2}{p}\right).$$

$\frac{n_1}{n} \xrightarrow{P} p$ and $\frac{n_0}{n} \xrightarrow{P} (1-p)$, with $p \in (0, 1)$.

The concern with the conventional estimator $\hat{\tau}$ is that it may be imprecise. In particular in settings where the outcome distribution is thick tailed, sometimes extremely so, confidence intervals may be wide. At the same time, it is possible that there is substantial evidence in the data that the treatment does have some effect. Randomization based tests on ranks or other transformations of the outcome (*e.g.*, Rosenbaum (2002)) may offer such evidence. However, the magnitude of the causal effect is often of primary interest, and it is not sufficient to simply reject the null of no effect. We address this by incorporating ideas from the robust estimation literature (Hampel et al., 2011; Huber, 2011) and by imposing some restrictions on the joint distribution of the potential outcomes motivated by the semiparametric literature (Bickel and Lehmann, 1975a,b, 1976, 2012) to develop new estimators.

The quantile treatment effect (Lehmann and D'Abrera, 1975) plays an important role in our setup. For quantile $u \in (0, 1)$, define

$$\Delta(u) = F_1^{-1}(u) - F_0^{-1}(u), \quad 0 \leq u \leq 1, \quad (2.3)$$

which is closely related to what Doksum (1974) and Doksum and Sievers (1976) label the *response* function: $\tilde{\Delta}(y) = F_1^{-1}(F_0(y)) - y = \Delta(F_0(y))$. The natural estimate for the quantile treatment effect $\Delta(u)$ is the empirical plug-in,

$$\hat{\Delta}(u) = \hat{F}_1^{-1}(u) - \hat{F}_0^{-1}(u)$$

where $\hat{F}_1^{-1}(u)$ equals $Y_{([n_1 u])}^{(1)}$, defined as the $[n_1 u]^{\text{th}}$ order statistic of $Y_i | Z_i = 1$, $i = 1, \dots, n_1$, where $n_1 = \sum_{i=1}^n Z_i$ and similarly for $\hat{F}_0^{-1}(u)$.

A natural class of parameters summarizing the difference between the $Y_i(1)$ and $Y_i(0)$ distributions consists of weighted averages of the quantile treatment effects:

$$\tau(F_0, F_1; W) = \int_0^1 \Delta(u) dW(u)$$

where the weights integrate to one, $W(0) = 0, W(1) = 1$. Different choices for the weight function correspond to different estimands. The constant weight case, $W'(u) \equiv 1$, corresponds to the population average treatment effect $\tau = \mathbb{E}(Y(1) - Y(0))$. The median corresponds to the case where W puts all its mass at $1/2$. We thus allow W to permit point masses.

For a given weight function W we can estimate the parameter $\tau(F_0, F_1; W)$ using a weighted average quantile estimator:

$$\begin{aligned} \hat{\tau}_W &\equiv \tau(\hat{F}_0, \hat{F}_1; W) = \int_0^1 \left(\hat{F}_1^{-1}(u) - \hat{F}_0^{-1}(u) \right) dW(u) \\ &= \frac{1}{n_1} \sum_{i=1}^{n_1} w_i^{(1)} Y_{(i)}^{(1)} - \frac{1}{n_0} \sum_{i=1}^{n_0} w_i^{(0)} Y_{(i)}^{(0)} + o_P(1) \end{aligned} \quad (2.4)$$

where

$$w_i^{(z)} = W\left(\frac{i}{n_z}\right) - W\left(\frac{i-1}{n_z}\right),$$

and $Y_{(i)}^{(z)}$ again are the order statistics in treatment group z .

2.2. Efficient estimation

We begin by considering the nonparametric model for a single sample, with cumulative distribution F where the interest is in the weighted quantile $\int_{-\infty}^{\infty} F^{-1}(u) dW(u)$ for a given weight function $W(\cdot)$. For this one-sample case, Jaeckel (1971a,b) builds on Chernoff et al. (1967) to show that under simple conditions on W and F , for a sample size of n ,

$$\int_0^1 \left(\hat{F}^{-1}(u) - F^{-1}(u) \right) dW(u) = \int_{-\infty}^{\infty} \psi(x, F) d(\hat{F}(x) - F(x)) + o_P(n^{-1/2}), \quad (2.5)$$

where the influence function ψ is related to the weight function W by

$$\psi(x, F) = - \int_x^{\infty} \frac{1}{f(y)} dW(F(y)) + \int_{-\infty}^{\infty} \frac{F(y)}{f(y)} dW(F(y)). \quad (2.6)$$

the last term ensures that $\int_{-\infty}^{\infty} \psi(x, F) dF(x) = 0$.

Note that if the derivatives ψ' and $W' = w$ exist, by (2.6),

$$\psi(x, F) = - \int_x^{\infty} w(F(y)) dy + \int_{-\infty}^{\infty} F(y) w(F(y)) dy$$

so that $\psi'(x, F) = w(F(x))$. Note that for the median, (2.6) yields, $\psi(x, F) = \frac{\text{sign}(x - F^{-1}(\frac{1}{2}))}{2f(F^{-1}(\frac{1}{2}))}$.

Our formula (2.6) is slightly more general than Jaeckel's, and in the appendix we establish sufficient conditions on F and W for our version of his result to hold.

[6] *Athey et al.*

Expression (2.5) in turn implies that

$$\int_0^1 \left(\hat{F}^{-1}(u) - F^{-1}(u) \right) dW(u) \xrightarrow{d} \mathcal{N}(0, \sigma^2(F, W)).$$

where the variance equals

$$\sigma^2(F, \omega) = \int_{-\infty}^{\infty} \psi(x, F)^2 dF(x).$$

Jaekel (1971a,b)'s results extend in the following way to the two-sample setting we are interested in. If $\tau(F_0, F_1, W)$ is estimated by

$$\hat{\tau}(F_0, F_1, W) = \int_0^1 (\hat{F}_1^{-1}(u) - \hat{F}_0^{-1}(u)) dW(u) \quad (2.7)$$

then, under regularity conditions given in the appendix (Theorem A.1),

$$\hat{\tau}(F_0, F_1, W) = \tau(F_0, F_1, W) + \frac{1}{n} \sum_{i=1}^n \left(Z_i \frac{\psi(Y_i, F_1)}{p} - (1 - Z_i) \frac{\psi(Y_i, F_0)}{1 - p} \right) + o_P(n^{-1/2}) \quad (2.8)$$

where $\psi(y, F)$ is given by (2.6) for given W .

Now let us return to the primary focus of this section, the estimation of the average treatment effect under the constant treatment effect assumption. In that case the quantile treatment effects are all equal:

$$\tau(u) \equiv \tau.$$

Then, for all weight functions W ,

$$\tau(F_0, F_1, W) = \tau. \quad (2.9)$$

Suppose that in addition to knowing that the treatment effect is constant, we know F_0 up to a shift. That is, $F_0(x) = F(x - \eta)$ where $F(\cdot)$ is known with derivative f , η unknown, and, because of the constant treatment effect, $F_1(x) = F(x - \eta - \tau)$. For this fully parametric model (with unknown parameters η and τ), if the Fisher information $I(f) = \int \left(\frac{f'}{f} \right)^2(x) f(x) dx = \int \left(-\frac{f'}{f} \right)'(x) f(x) dx$ satisfies $0 < I(f) < \infty$, the maximum likelihood estimator of τ , suitably regularized (*e.g.*, Le Cam and Yang (1988)), has influence function,

$$\psi_{f,\eta}(Z, Y; \tau) = -\frac{1}{I(f)} \cdot \left(\frac{Z}{p} \cdot \frac{f'}{f}(Y - \eta - \tau) - \frac{1 - Z}{1 - p} \cdot \frac{f'}{f}(Y - \eta) \right). \quad (2.10)$$

If f'/f is absolutely continuous, then

$$w_f(F(x)) = \frac{1}{I(f)} \left(-\frac{f'}{f} \right)'(x) \text{ or } w_f(u) = \frac{1}{I(f)} \left(-\frac{f'}{f} \right)'(F^{-1}(u)) \quad (2.11)$$

provides an efficient L estimate when substituted appropriately in (2.7), leading to

$$\hat{\tau}(F_0, F_1, W) = \int_0^1 (\hat{F}_1^{-1}(u) - \hat{F}_0^{-1}(u)) w_f(u) du. \quad (2.12)$$

2.3. Fully adaptive estimation

It follows from semiparametric theory (Bickel et al., 1993) that even if the density f is unknown, substituting a suitable estimate of f (and η) in (2.10) or (2.12), will yield estimators with influence functions given by (2.10), or equivalently by

$$\psi_{f_0}(Z, Y; \tau) = -\frac{1}{I(f_0)} \cdot \left(\frac{Z}{p} \cdot \frac{f'_0}{f_0}(Y - \tau) - \frac{1-Z}{1-p} \cdot \frac{f'_0}{f_0}(Y) \right), \quad (2.13)$$

where $f_0(\cdot) \equiv F'_0(\cdot) \equiv f(\cdot - \eta)$.

This is a consequence of the orthogonality of the tangent space with respect to f .

Formally we have for the unknown $f(\cdot)$ case:

Theorem 1. *For all f such that f' exists and $0 < I(f) < \infty$:*

- (i) *There exist $\sqrt{n}$ consistent estimators $\hat{\tau}$ and $\hat{\mu}$.*
- (ii) *Under mild conditions (see (A.9) and (A.10) in the appendix), we can construct an M estimate $\hat{\tau}^{if}$ such that*

$$\sqrt{n}(\hat{\tau}^{if} - \tau) \xrightarrow{d} \mathcal{N}\left(0, \frac{1}{p(1-p)I(f)}\right). \quad (2.14)$$

- (iii) *Under mild conditions (see Lemma A.2 in the appendix), we can construct an L estimate $\hat{\tau}^{waq}$ (weighted average quantile) such that $\hat{\tau}^{waq}$ satisfies (2.14).*

Our conditions are not optimal (see Stone (1975) for minimal ones in the one sample case).

The M estimate is of the form,

$$\hat{\tau}^{if} \equiv \tilde{\tau} + \frac{1}{n} \left\{ \sum_{i=1}^{n/2} \psi_{\hat{f}_{0(2)}}(Z_i, Y_i; \tilde{\tau}) + \sum_{i=1+n/2}^n \psi_{\hat{f}_{0(1)}}(Z_i, Y_i; \tilde{\tau}) \right\} \quad (2.15)$$

where $\hat{f}_{0(1)}$ is an estimate of f_0 using $\{(Z_i, Y_i) : 1 \leq i \leq \frac{n}{2}\}$, and $\hat{f}_{0(2)}$ using $\{(Z_i, Y_i) : \frac{n}{2} < i \leq n\}$, a one step estimate using the sample splitting technique (Klaassen, 1987).

The L estimate is obtained by direct plug-in of weight function estimation using sample splitting as above. Details are in the appendix.

Thus, both the estimators $\hat{\tau}^{waq}$ and $\hat{\tau}^{if}$ are adaptive to models in which the distribution of control and treated potential outcomes is known up to location. Theorem 1 implies that the influence function based estimator is as efficient as the maximum likelihood estimator based on the true distribution function. For instance, if potential

outcomes are normally distributed, the maximum likelihood estimator is the difference in means, and the influence function estimator has the same limiting distribution. If, however, the potential outcomes follow a double exponential distribution, the difference in means estimator is not efficient; the difference in medians is the efficient estimator. Under this true distribution, the influence function based estimator adapts and has the same limiting distribution as the difference in medians. For the Cauchy distribution the optimal weights are more complicated, $w_f(u) \propto -\cos(2\pi u)\sin(\pi u)^2$, but the influence function based estimator has the same limiting distribution, without requiring a priori knowledge about the distribution.

This influence function based estimator $\hat{\tau}^{if}$ is a special case of the estimator developed by Cuzick (1992a,b) for the partial linear regression model setting.

Although these estimators are efficient under the constant additive treatment model, as we shall see in Section 3, if the model is violated, estimates of the types derived from f known continue to estimate at rate $n^{-1/2}$, meaningful measures of the treatment effect as discussed in Section 2.1. This is unfortunately not the case for $\hat{\tau}^{if}$ and $\hat{\tau}^{waq}$ because estimation of f' and f'' introduces components of variance of order larger than $n^{-1/2}$. However, there is a partial remedy.

2.4. Partial Adaptation

An interesting alternative to full adaptation in this problem was studied by Jaeckel (1971b) in the one-sample symmetric case. See also Yu and Yao (2017). The estimator proposed by Jaeckel (1971b) is

$$\hat{\mu}_\alpha \equiv \frac{1}{1-2\alpha} \int_\alpha^{1-\alpha} \hat{F}^{-1}(u) du$$

for a sample from $F(x - \mu)$ with f symmetric. We can interpret this as restricting the class of weight functions to one indexed by a scalar parameter α :

$$w_\alpha(u) = \begin{cases} \frac{1}{1-2\alpha} & \text{if } \alpha \leq u \leq 1 - \alpha \\ 0 & \text{otherwise.} \end{cases}$$

This weight function yields the α -trimmed mean. We can then choose the value of α that minimizes the asymptotic variance of $\hat{\mu}_\alpha$. This asymptotic variance is equal to

$$\begin{aligned} \sigma_\alpha^2 \equiv & \frac{1}{n(1-2\alpha)^2} (\mathbb{E}(X - \bar{\mu})^2 1(F^{-1}(\alpha) \leq X \leq F^{-1}(1-\alpha)) \\ & + \alpha \cdot (F^{-1}(1-\alpha) - \bar{\mu})^2 + \alpha(F^{-1}(\alpha) - \bar{\mu})^2), \end{aligned}$$

with

$$\bar{\mu} = \int_\alpha^{1-\alpha} F^{-1}(u) du + \alpha(F^{-1}(\alpha) + F^{-1}(1-\alpha)),$$

which can be estimated by replacing F with the empirical distribution, denoted as $\hat{\sigma}_\alpha^2$. Let $\hat{\alpha} = \arg \min_\alpha \hat{\sigma}_\alpha^2$, and let $\hat{\mu}_{\hat{\alpha}}$ be the corresponding estimator for μ . Jaeckel (1971b) shows that $\hat{\mu}_{\hat{\alpha}}$ is adaptive for estimating μ over a Huber family of densities. In the Huber family $(X - \mu)/\sigma$ has density f for varying $\mu, \sigma > 0$:

$$\log f(x) = \begin{cases} -\frac{x^2}{2} - c(k), & \text{if } |x| \leq k \\ -\frac{k|x|}{2} - c(k) & \text{if } |x| > k. \end{cases}$$

where $c(k)$ makes $\int f(x)dx = 1$, and $k = -F^{-1}(\alpha)$. Adaptivity here means that using the trimming proportion optimizing the variance estimate, in fact yields an estimate which is efficient for the member of the Huber family generating the data. He optimizes $0 < \alpha_0 \leq \alpha \leq \alpha_1 < \frac{1}{2}$.

Because this family includes among others the Gaussian ($k \rightarrow \infty$) and double exponential ($k = 0$), this family is very flexible. For more properties, see Huber (2011).

In the two-sample case, it is reasonable to consider asymmetric weight functions leading to the natural generalization,

$$\hat{\tau}_{\alpha,\beta} = \frac{1}{1 - (\alpha + \beta)} \int_\alpha^{1-\beta} (\hat{F}_1^{-1}(u) - \hat{F}_0^{-1}(u)) du.$$

We show in the online appendix that Jaeckel's result on partial adaptation can, for our two sample problem, be extended to a generalization of the Huber family whose members are symmetric iff $k_1 = k_2$, defined by $(X - \mu)/\sigma \sim f$:

$$\log f(x) = \begin{cases} -\frac{x^2}{2} - c(k_1, k_2), & \text{if } -k_1 \leq x \leq k_2, \\ \frac{k_1 x}{2} - c(k_1, k_2), & \text{if } x < -k_1, \\ -\frac{k_2 x}{2} - c(k_1, k_2), & \text{if } x > k_2, \end{cases} \quad (2.16)$$

where $c(k_1, k_2) \equiv \log\left(2 \left(\frac{e^{-k_1/2}}{k_1} + \frac{e^{-k_2/2}}{k_2}\right) + \sqrt{2\pi}(\Phi(k_2) - \Phi(-k_1))\right)$ and Φ is the CDF of $\mathcal{N}(0, 1)$.

This family can be equivalently parametrized by $F(-k_1)$ and $F(k_2)$.

3. The General Parametric Treatment Effect Case

In some settings, the assumption of an additive model may be too restrictive. In this section, we develop estimators given a general parametric model for this difference.

The starting point is a model governing the relation between the two potential outcomes:

Assumption 1 (Parametric Model for Unit-level Treatment Effects). *The unit-level potential outcomes satisfy*

$$Y_i(1) = h(Y_i(0), \theta), \quad \text{with } \frac{\partial}{\partial y} h(y, \theta) > 0,$$

[10] *Athey et al.*

for some known function h and an unknown parameter, $\theta \in \mathcal{R}$. Equivalently,

$$F_1(h(y, \theta)) = F_0(y).$$

As before we initially assume F_0 known. In terms of the quantile treatment effects $\tau(u)$ Assumption 1 implies the restriction

$$\tau(u) = F_1^{-1}(u) - F_0^{-1}(u) = h(F_0^{-1}(u), \theta) - F_0^{-1}(u).$$

Given Assumption 1, the population average treatment effect can be characterized as

$$\tau^{\text{pop}} = \int_0^1 (h(F_0^{-1}(u), \theta) - F_0^{-1}(u)) du$$

In practice, however, estimating τ^{pop} may still be subject to substantial sampling variance. As an alternative, we suggest to estimate the in-sample, as opposed to population, average treatment effect. An insight is that estimators of this object can have a much lower variance. We define the in-sample average treatment effect as

$$\begin{aligned} \tau^{\text{is}} &= \frac{1}{N} \sum_{i=1}^N (Y_i(1) - Y_i(0)) \\ &= \frac{1}{N} \sum_{i=1}^N \left\{ Z_i(Y_i - h^{-1}(Y_i, \theta)) + (1 - Z_i)(h(Y_i, \theta) - Y_i) \right\}. \end{aligned} \tag{3.17}$$

When h is not just an additive function, τ^{is} is sample-dependent, and thus stochastic. In particular when the variance of Y_i is large because of thick tails for the potential outcome distributions, the variance of τ^{is} will be large, too. We therefore focus on the variance of estimators $\hat{\tau}$ relative to τ^{is} for the particular sample at hand, rather than relative to the population average τ^{pop} .

Because F_0 is known, we can estimate θ efficiently by some version of maximum likelihood to get an estimate $\hat{\theta}$ and

$$\hat{\tau} = \int_0^1 (h(F_0^{-1}(y), \hat{\theta}) - F_0^{-1}(y)) dy$$

as an estimate of τ . The density of Y given Z is

$$f(y|z) = (f_0(y))^{1-z} (f_1(y, \theta))^z.$$

By Assumption 1

$$f_1(y, \theta) = f_0(h^{-1}(y, \theta)) \frac{\partial h^{-1}(y, \theta)}{\partial y}, \tag{3.18}$$

and the score function is

$$\dot{\ell}(y, z, \theta) \equiv z \cdot \frac{\partial}{\partial \theta} \log f_1(y, \theta)$$

yielding,

$$\hat{\theta} = \theta + \frac{1}{n} \sum_{i=1}^n I^{-1} \dot{\ell}(Y_i, Z_i, \theta) + o_P(n^{-1/2})$$

where $I = E \left(\dot{\ell}(Y, Z, \theta) \right)^2$.

If F_0 is assumed unknown, to obtain an efficient influence function we must

- 1) Compute the tangent plane as f_0 varies with θ fixed. The tangent plane is given by

$$\dot{P}_f = \{u(Y, Z) = (1-Z)v(Y, \theta) + Zv(h^{-1}(Y, \theta), \theta) : \int v^2(y, \theta) f_0(y) dy < \infty, \int v(y, \theta) f_0(y) dy = 0\}.$$

(Note that both factors of the likelihood must be varied treating θ as fixed.)

- 2) Project $\dot{\ell}$ on the orthocomplement of the tangent plane to get

$$\dot{\ell}^*(Y, Z, \theta) = \dot{\ell}(Y, Z, \theta) - (ZQ(h^{-1}(Y, \theta), \theta) + (1-Z)Q(Y, \theta))$$

where $ZQ(h^{-1}(Y, \theta), \theta) + (1-Z)Q(Y, \theta)$ is the projection of $\dot{\ell}$ on $\dot{P}_f$.

- 3) The efficient influence function is given by

$$\psi(Y, Z, \theta) = \frac{\dot{\ell}^*(Y, Z, \theta)}{E \left(\dot{\ell}^*(Y, Z, \theta) \right)^2}.$$

Lemma 1. *The efficient influence function for θ is*

$$\psi_{f_0}(y, z, \theta) = \mathcal{I}^{-1} \left\{ \frac{z}{p} \cdot g(y, \theta) - \frac{1-z}{1-p} \cdot g(h(y, \theta), \theta) \right\},$$

where

$$g(y, \theta) = \frac{\partial}{\partial \theta} \log \left(f_0(h^{-1}(y, \theta)) \cdot \frac{\partial h^{-1}(y, \theta)}{\partial y} \right)$$

and

$$\mathcal{I} = \int g^2(h(y, \theta), \theta) f_0(y) dy.$$

To use the influence function approach we need a $\sqrt{n}$ -consistent initial estimator $\tilde{\theta}$. We can do so by a fixed number of quantiles, $u_1, \dots, u_d$, where d is the dimension of θ , and find the θ that solves

$$\hat{F}_1^{-1}(u) = h(\hat{F}_0^{-1}(u), \theta),$$

for $u = u_1, \dots, u_d$. Let $\tilde{\theta}$ denote the solution to this system of equations. Then:

[12] *Athey et al.*

Theorem 2. *Under mild conditions on the estimation of density and its derivative, the estimator $\hat{\theta}^{if}$ below is efficient for θ , i.e. $\sqrt{n}(\hat{\theta} - \theta) \xrightarrow{d} \mathcal{N}(0, \frac{1}{p(1-p)}\mathcal{I}^{-1})$:*

$$\hat{\theta}^{if} \equiv \tilde{\theta} + \frac{1}{n} \left\{ \sum_{i=1}^{n/2} \psi_{\hat{f}_{0(2)}}(Y_i, Z_i; \tilde{\theta}) + \sum_{i=1+n/2}^n \psi_{\hat{f}_{0(1)}}(Y_i, Z_i; \tilde{\theta}) \right\}$$

where $\hat{f}_{0(1)}$ is the estimate of f_0 using $\{(Y_i, Z_i) : 1 \leq i \leq \frac{n}{2}\}$, and $\hat{f}_{0(2)}$ is the estimate of f_0 using $\{(Y_i, Z_i) : \frac{n}{2} < i \leq n\}$, again a one step estimate using the sample splitting technique (Klaassen, 1987), similar to (2.15).

Note that, just as in the constant treatment effect setting, the efficient influence function is not the one corresponding to F_0 known.

4. Simulations

We evaluate the performance of the proposed estimators and conventional estimators in a Monte Carlo study. Throughout most of these simulations, the true unit-level treatment effects are all zero. We estimate the treatment effect using the proposed efficient estimators based on an additive model. We consider seven estimators: The (standard) difference in means, the difference in medians, the Hodges-Lehman (Hodges Jr and Lehmann (1963)) estimator,[‡] the adaptively trimmed mean, the adaptively winsorized mean,[§] the estimator based on the efficient influence function (eif), and the weighted average quantiles (waq) estimator. For the latter two we report only results without sample splitting. Results for the case where we do use sample splitting are very similar and are available from the authors.

We present results for three sets of simulations. First, we report results for settings with a wide range of known distributions for the potential outcomes, so we can directly assess the ability of the proposed methods to adapt to different distributions. Second, we report results for simulations based on real data, one set based on housing prices and one based on medical expenditures. Both data sets are examples of relatively heavy tailed distributions as commonly encountered in empirical work.

4.1. Simulations with Known Distributions

We simulate samples of $n = 20,000$ observations, exactly half of which are treated, and report summary statistics based on 10,001 simulated samples (using an odd number so

[‡]The Hodges-Lehmann estimator is equal to the median of all pairwise differences between treated and control observations.

[§]For both trimming and winsorizing, we choose the trimming percentage that minimizes an estimate of the asymptotic variance, allowing anywhere between no trimming (difference in means) and the extreme of trimming all but the medians (difference in medians). While including the extremes is not covered by the theory, this approach appears to work well in our simulations.

Fig. 1. The efficient weight function for the Cauchy distribution. The weights, normalized to have mean 1, are plotted against the quantile (left) and against the value of the observations (right). For the figure on the right, we only plot the range from -10 to 10 , which corresponds to approximately the 0.03 to 0.97 quantile. At this point, the weight is approximately -0.04 , and weights for more extreme quantiles are closer to 0.

that the median is unique). We repeat the simulation study for standardized Normal, Double Exponential (Laplace), and Cauchy distributions for the potential outcomes. The difference in means is the maximum likelihood estimator Normally distributed data, and so will do well there, but may perform poorly for thicker tailed distributions such as the Double Exponential distribution and in particular the Cauchy distribution. The difference in medians is the maximum likelihood estimator for the Double Exponential distribution, and relatively robust to thick tails and outliers, and so is expected to perform reasonably well across all specifications, but not as well as the efficient estimators for the Normal.

For the simulations with known distributions we can infer the functional form for the optimal weights for the quantile-based estimator. The optimal weights for the waq estimator are proportional to the (estimated) second derivative of the log density. For the Normal distribution, $\frac{\partial^2 \ln f}{\partial y^2}(y) = -\frac{1}{\sigma^2}$, implying the optimal weights are constant, so that the semiparametrically efficient estimator approximates the simple difference in means (which is the parametric maximum likelihood estimator in this case). The density of the double exponential distribution is such that the optimal weights asymptotically place all weight close to the median. The maximum likelihood estimator for the location of a Double Exponential distribution is the sample median. For the standard Cauchy distribution the efficient weights w on the difference in $u \in (0, 1)$ quantiles of treated and control distributions are $w_f(u) \propto -\cos(2\pi u) \sin(\pi u)^2$, shown in Figure 1. Most of the weight is concentrated around the median, with strictly negative weights outside the $[0.25, 0.75]$ quantile range.

The efficient estimators perform well across distributions, and confidence intervals

Table 1. Summary statistics for simulations with different distributions: Normal, Double Exponential, Cauchy. Statistics shown are based on 10,001 simulated samples. Each sample has 10,000 treated observations and 10,000 control observations. Relative efficiency is the ratio of standard deviation to the square root of the efficiency bound. Columns labeled length show the median length of the confidence intervals.

estimator	bias	standard deviation	relative efficiency	RMSE	MAD	95% C.I.		boot. var. C.I.	
						coverage	median length	coverage	median length
Normal Distribution:									
diff. in means	0.000	0.014	1.01	0.014	0.010	0.95	0.055	0.95	0.055
diff. in medians	-0.000	0.018	1.26	0.018	0.012	0.95	0.069	0.95	0.069
Hodges-Lehmann	0.000	0.015	1.03	0.015	0.010	0.95	0.057	0.95	0.057
adaptive trim	0.000	0.015	1.03	0.015	0.010	0.94	0.055	0.95	0.057
adaptive wins.	0.000	0.014	1.01	0.014	0.010	0.95	0.055	0.95	0.055
eif	0.000	0.014	1.02	0.014	0.010	0.95	0.056	0.95	0.057
waq	0.000	0.014	1.02	0.014	0.010	0.95	0.056	0.95	0.056
Double Exponential Distribution:									
diff. in means	-0.000	0.020	1.43	0.020	0.014	0.95	0.078	0.95	0.078
diff. in medians	0.000	0.014	1.01	0.014	0.010	0.97	0.061	0.95	0.057
Hodges-Lehmann	-0.000	0.016	1.17	0.016	0.011	0.95	0.064	0.95	0.064
adaptive trim	0.000	0.014	1.02	0.014	0.010	0.95	0.059	0.95	0.057
adaptive wins.	-0.000	0.018	1.29	0.018	0.012	0.96	0.076	0.95	0.069
eif	-0.000	0.015	1.06	0.015	0.010	0.95	0.060	0.96	0.060
waq	-0.000	0.015	1.07	0.015	0.010	0.95	0.060	0.96	0.061
Cauchy Distribution:									
diff. in means	0.462	127.149	6357.47	127.144	2.047	0.98	9.324	0.98	9.292
diff. in medians	0.000	0.022	1.11	0.022	0.015	0.97	0.093	0.95	0.087
Hodges-Lehmann	0.000	0.026	1.28	0.026	0.018	0.95	0.101	0.95	0.101
adaptive trim	0.000	0.022	1.08	0.022	0.015	0.95	0.092	0.95	0.086
adaptive wins.	0.000	0.024	1.19	0.024	0.016	0.98	0.111	0.97	0.100
eif	0.000	0.020	1.01	0.020	0.014	0.97	0.085	0.97	0.088
waq	0.000	0.021	1.05	0.021	0.014	0.96	0.085	0.98	0.099

based either on estimates of the analytic variance formulas or on the bootstrap achieve their nominal coverage levels, as shown in Table 1. Their standard deviations are close to the theoretical efficiency bound, as shown in column 4, labelled relative efficiency, where values larger than one imply standard deviations of the estimator in excess of the efficiency bound. The efficient influence function and weighted average quantiles estimators are close to the most efficient estimator for the normal, double exponential, and Cauchy distribution. The last columns show that the confidence intervals are close to their nominal coverage for each distribution. For computational convenience in the simulations, the confidence intervals based on the bootstrap variance use the m -out-of- n bootstrap (Bickel et al. (2012)), with $m = 2,000$ (half treated, half control), to estimate the variance of the estimators. Even with these smaller sample sizes, the density estimates calculated in the bootstrap samples appear to be sufficiently good to yield reasonable confidence intervals for the estimators.

Fig. 2. Histogram of house prices, in levels and logs, as well as estimated optimal weights (second derivative of the log density) based on all 170,239 observations. The weights are normalized to be mean 1 (black horizontal line). Some weights are below 0 (blue horizontal line). Vertical lines indicate the 0.0001, 0.001, 0.01, 0.1, 0.9, 0.99, 0.999, 0.9999 quantiles.

4.2. Simulations with House Price Data

In the second set of simulations we use house price data from the replication files of Linden and Rockoff (2008, 2019). They obtained property sales data for Mecklenburg County, North Carolina, between January 1994 and December 2004. They dropped sales below \$5,000 and above \$1,000,000, such that 170,239 observations remain, which we take as our population of interest. Despite the trimming the distribution is noticeably skewed (skewness 2.2) and heavy-tailed (kurtosis 9.5). Even after taking logs, the distribution is heavy-tailed with kurtosis equal to 5.1. Figure 2 plots a histogram for house prices, both in levels and in logs, along with the estimated optimal weights (second derivative of the log density) based on all 170,239 observations.

We base simulations on this data by drawing samples of size $n = 20,000$, and randomly assigning exactly half of each sample to the treatment group and the remaining half to the control group, with a zero treatment effect. Each sample is drawn without replacement, but observations may appear in multiple samples. Because the true distribution, however, is not available analytically, we cannot calculate the efficiency bound $\mathcal{I}$ directly. Instead, we estimate the efficiency bound using density estimates based on all observations. We compute the same estimators as in the simulations of the previous section, with a small adjustment to the adaptively trimmed and winsorized means. In settings with positively-valued outcomes, researchers are typically more concerned about the effect of heavy right, rather than left, tails on the variance of estimators. We therefore fix the trimming and winsorizing percentiles on the left to 0% (no trimming), and only adaptively choose the threshold on the right. Without such an adjustment, the estimators tend to trim (too) little in order to avoid losing observations from the left tail, which are unproblematic for the precision of the estimator.

Table 2 summarizes the simulation results based on 10,001 simulated samples. When the house prices are in levels, the standard deviation of the difference in means estimator is twice as large as that of efficient estimators. For the difference in medians and the Hodges-Lehman estimators, which are less affected by outliers in the data, the standard deviation is larger by approximately 30% and 20%, respectively. Confidence intervals, based on estimated variances and asymptotic normal approximations, have close to nominal coverage throughout, and are meaningfully shorter for the efficient estimators we propose.

We also estimate a proportional treatment effect (multiplicative) model and translate the estimated coefficients into level effects. Under the multiplicative model, $Y_i(1) = \theta Y_i(0)$. When outcomes are strictly positive, this is identical to an additive model for log outcomes; $\log(Y_i(1)) = \log(\theta) + \log(Y_i(0))$. Given an estimate $\hat{\tau}_{\log}$ of the additive model with the outcome in logs, the estimate of the level treatment effect is then

$$\hat{\tau} = \frac{1}{n} \left(\sum_{i=1}^n (1 - Z_i) ((\exp(\hat{\tau}_{\log}) - 1) Y_i) + Z_i ((1 - \exp(-\hat{\tau}_{\log})) Y_i) \right).$$

For estimates of the in-sample treatment effect, we therefore apply estimates $\hat{\tau}_{\log}$ to a population of interest with known means μ_{Y_0} and μ_{Y_1} and fixed treatment probability p as

$$\hat{\tau} = (1 - p) ((\exp(\hat{\tau}_{\log}) - 1) \mu_{Y_0}) + p ((1 - \exp(-\hat{\tau}_{\log})) \mu_{Y_1}).$$

Using the Delta method, if V is the asymptotic variance of $\hat{\tau}_{\log}$, then the asymptotic variance of $\hat{\tau}$, holding μ_{Y_0} , μ_{Y_1} , and p fixed, is

$$\left((1 - p) \exp(\tau_{\log}) \mu_{Y_0} + p \exp(-\tau_{\log}) \mu_{Y_1} \right)^2 V$$

which we estimate by replacing $\tau_{\log}$ with $\hat{\tau}_{\log}$ and V by the estimate of the variance, $\hat{V}$. For the purpose of these simulations, we set $\mu_{Y_0} = \mu_{Y_1}$ equal to the population mean of house prices, and $p = 1/2$.

When treatment effects are assumed to be proportional to potential outcomes, the proposed estimators for this multiplicative model are still more efficient than alternative estimators, but the gains are smaller. The middle panel of Table 2 shows the quality of estimates of the multiplicative parameter obtained by transforming outcomes into logs, $\log(\theta)$. As can be seen in panel (b) of Figure 2, the distribution of log house prices appears closer to the normal distribution with fewer “outliers.” Consequently, the difference in means estimator, which is efficient for the normal distribution, comes noticeably closer to the efficiency bound than when outcomes are in levels. Nevertheless, the efficient estimators further reduce the variance. The bottom panel of Table 2 shows the same summary statistics when the multiplicative parameter is translated back into a level effect. The efficient estimators of the multiplicative parameter lead to treatment effects

Table 2. Summary statistics for simulations based on house price data from Linden and Rock-off (2008). Statistics shown are based on 10,001 simulated samples. Each sample has 10,000 treated observations and 10,000 control observations. Relative efficiency is the ratio of standard deviation to the square root of the efficiency bound, which we calculate from density estimates based on the full sample. Columns labeled length show the median length of the confidence intervals. The adaptively trimmed mean and adaptively winsorized mean only apply trimming at the top, but not at the bottom.

estimator	bias	standard deviation	relative efficiency	RMSE	MAD	95% C.I.		boot. var. C.I.	
						coverage	length	coverage	length
effect in levels based on additive model in levels									
diff. in means	-28	1871	1.92	1871	1251	0.95	7343	0.95	7339
diff. in medians	8	1259	1.30	1259	855	0.95	4989	0.95	4893
Hodges-Lehmann	-5	1138	1.17	1138	757	0.95	4487	0.95	4487
adaptive trim	5	989	1.02	989	651	0.98	4560	0.95	3914
adaptive wins.	4	1114	1.15	1114	738	0.99	6378	0.97	4760
eif	13	941	0.97	941	628	0.95	3819	0.95	3768
waq	13	946	0.97	946	626	0.95	3819	0.95	3859
multiplicative parameter: additive model in logs									
diff. in means	-0.0000	0.0083	1.35	0.0083	0.0055	0.95	0.033	0.95	0.032
diff. in medians	0.0000	0.0075	1.23	0.0075	0.0051	0.95	0.030	0.95	0.029
Hodges-Lehmann	-0.0000	0.0072	1.17	0.0072	0.0048	0.95	0.028	0.95	0.028
adaptive trim	-0.0000	0.0083	1.35	0.0083	0.0055	0.95	0.033	0.95	0.032
adaptive wins.	-0.0000	0.0083	1.35	0.0083	0.0055	0.95	0.033	0.95	0.032
eif	-0.0000	0.0067	1.08	0.0067	0.0045	0.95	0.026	0.95	0.026
waq	-0.0000	0.0067	1.08	0.0067	0.0044	0.95	0.026	0.95	0.027
effect in levels based on additive model in logs									
diff. in means	-7	1695	1.35	1695	1125	0.95	6658	0.95	6657
diff. in medians	9	1545	1.23	1545	1048	0.95	6128	0.95	6001
Hodges-Lehmann	-6	1468	1.17	1468	976	0.95	5786	0.95	5783
adaptive trim	-7	1695	1.35	1695	1125	0.95	6658	0.95	6657
adaptive wins.	-7	1695	1.35	1695	1125	0.95	6658	0.95	6657
eif	-5	1363	1.08	1363	914	0.95	5374	0.95	5415
waq	-3	1364	1.08	1364	910	0.95	5374	0.95	5462

with smaller variance and shorter confidence intervals than the difference in means, irrespective of whether the latter is estimated in levels (top panel) or in logs and then translated into levels (bottom panel).

4.3. Medical Expenditures Data

Next, we present simulation results based on medical expenditure data from the IBM MarketScan Research Database, following the sample construction of Koenecke et al. (2021, Figure 2). We restrict the sample to males, age 45–64, with pneumonia inpatient diagnosis and at least 1 year of continuous medical enrollment. For each patient, we consider the first inpatient admission only to abstract away from any dynamics. We focus on medical expenditure as the outcome variable. For each patient, we sum the payments recorded by MarketScan for this admission. In total, we use data on 103,662

Fig. 3. Histogram of medical expenditures per admission, in levels and logs, as well as estimated optimal weights (second derivative of the log density), based on the 98,155 observations in the control group. For the figure in levels, vertical lines indicate the 0.001, 0.01, 0.1, and 0.9 quantiles; the figure is limited to below \$200,000, such that the 0.99 and higher quantiles do not appear. For the figure in logs, vertical lines indicate the 0.001, 0.01, 0.1, 0.99, 0.999, 0.9999 quantiles. The weights are normalized to be mean 1 (black horizontal line). Some weights are below 0 (blue horizontal line).

admissions.[¶] Figure 3 plots a histogram for medical expenditure in levels and in logs along with the estimated optimal weights (second derivative of the log density). We also observe a treatment variable in this data set, the (prior) use of alpha blockers, which Koenecke et al. (2021) find may improve health outcomes during respiratory distress by preventing hyperinflammation.

We design a simulation study similar to those of the previous sections, treating the receipt of alpha blockers as randomly assigned in our population. This allows us to study (coverage) properties of the estimators and inference procedures in settings where the parametric treatment effect model is not (necessarily) correctly specified. For these simulations, each sample is a draw, without replacement, of 200 of the 5,507 observations in the treatment group and 3,565 of the 98,155 observations in the control group. While the control group is smaller than in our other simulations, it remains sufficiently large to estimate the densities and its derivatives required for our estimators. In this simulation design, F_0 is given by the empirical distribution of the control group, and F_1 is given by the empirical distribution of the treatment group, of the full sample. Although it is not necessarily correct, the additive treatment effect model may offer a reasonable approximation, and we are interested in the performance of inference methods when the conditions for our theoretical results are not (quite) met in this particular application.

The simulation results reported in Table 3 suggest that the proposed estimators perform reasonably well in this setting despite mis-specification. The bias, RMSE, MAD,

[¶]The sample size differs slightly from that reported by Koenecke et al. (2021) due to missing expenditure data for a small number of admissions.

and coverage are relative to the population difference in means. The third panel of the table translate the multiplicative parameter into levels effects as in the previous section for simulations based on house prices. Here, we take $\mu_{Y_0} = 38,745$ and $\mu_{Y_1} = 34,872$, as well as $p = 0.0531$, as in the distribution these simulations draw from. As for the house price data, the adaptively trimmed mean and adaptively winsorized mean are asymmetric, only trimming the right tail of the data.

5. Conclusion

In many modern settings where randomized experiments are used to estimate treatment effects the presence of heavy-tailed distributions can lead to larger standard errors. Often researchers use Winsorizing with *ad hoc* thresholds to address this. Here we develop systematic methods for obtaining more precise inferences using parametric models for the treatment effects, while avoiding the specification of models for the potential outcomes. We present results for semiparametric efficiency bounds, suggest efficient estimators, and show in simulations that these methods can be effective in realistic settings.

Table 3. Summary statistics for simulations based on the medical expenditures data in logs, across 1,000 simulated samples. Each sample has 200 treated observations and 3,565 control observations. Coverage and median length refer to 95% confidence intervals.

estimator	bias	standard deviation	RMSE	MAD	95% C.I.		boot. var. C.I.	
					coverage	length	coverage	length
effect in levels based on additive model in levels								
error and coverage relative to population difference in means								
diff. in means	-108	4567	4566	3016	0.93	16935	0.93	16888
diff. in medians	3345	1271	3578	3222	0.46	6161	0.28	5259
Hodges-Lehmann	3490	891	3602	3464	0.03	3664	0.01	3534
adaptive trim	3877	641	3929	3860	0.00	3052	0.00	2778
adaptive wins.	2816	1384	3138	2786	0.52	5668	0.57	6023
eif	3928	682	3987	3922	0.01	4878	0.00	3240
waq	3846	792	3927	3848	0.03	4878	0.01	4018
multiplicative parameter: additive model in logs								
multiplicative parameter: error and coverage relative to population difference in means								
diff. in means	-0.001	0.077	0.077	0.052	0.95	0.30	0.95	0.30
diff. in medians	0.005	0.080	0.080	0.051	0.97	0.33	0.95	0.32
Hodges-Lehmann	0.008	0.071	0.071	0.048	0.95	0.29	0.95	0.28
adaptive trim	0.016	0.074	0.076	0.053	0.95	0.29	0.95	0.29
adaptive wins.	-0.002	0.078	0.078	0.052	0.94	0.30	0.95	0.31
eif	0.029	0.066	0.072	0.049	0.95	0.27	0.94	0.26
waq	0.028	0.066	0.071	0.047	0.95	0.27	0.95	0.27
effect in levels based on additive model in logs								
effect in levels: error and coverage relative to population difference in means								
diff. in means	2423	2871	3756	2522	0.90	11167	0.91	11138
diff. in medians	2647	3009	4006	2499	0.92	12318	0.92	11951
Hodges-Lehmann	2747	2673	3832	2672	0.88	10730	0.87	10390
adaptive trim	3062	2809	4154	2983	0.86	11055	0.85	10850
adaptive wins.	2365	2915	3752	2544	0.90	11084	0.92	11389
eif	3532	2525	4341	3481	0.78	10452	0.76	10014
waq	3464	2523	4285	3413	0.79	10441	0.77	10147

References

- Batir, N. (2017). Bounds for the gamma function. *Results in Mathematics* 72, 865–874.
- Bickel, P. J. (1967). Some contributions to the theory of order statistics. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pp. 575–591. University of California Press.
- Bickel, P. J. (1982). On adaptive estimation. *The Annals of Statistics* 10(3), 647–671.
- Bickel, P. J. and K. A. Doksum (2015). *Mathematical statistics: basic ideas and selected topics*, Volume 2. CRC Press.
- Bickel, P. J., F. Götze, and W. R. van Zwet (2012). Resampling fewer than n observations: gains, losses, and remedies for losses. In *Selected works of Willem van Zwet*, pp. 267–297. Springer.
- Bickel, P. J., C. A. Klaassen, Y. Ritov, and J. A. Wellner (1993). *Efficient and adaptive estimation for semiparametric models*, Volume 4. Johns Hopkins University Press Baltimore.
- Bickel, P. J. and E. L. Lehmann (1975a). Descriptive statistics for nonparametric models i. introduction. *The Annals of Statistics* 3(5), 1038–1044.
- Bickel, P. J. and E. L. Lehmann (1975b). Descriptive statistics for nonparametric models ii. location. *The Annals of Statistics* 3(5), 1045–1069.
- Bickel, P. J. and E. L. Lehmann (1976). Descriptive statistics for nonparametric models. iii. dispersion. *Annals of Statistics* 4(6), 1139–1158.
- Bickel, P. J. and E. L. Lehmann (2012). Descriptive statistics for nonparametric models iv. spread. In *Selected Works of EL Lehmann*, pp. 519–526. Springer.
- Chernoff, H., J. L. Gastwirth, and M. V. Johns (1967). Asymptotic distribution of linear combinations of functions of order statistics with applications to estimation. *The Annals of Mathematical Statistics* 38(1), 52–72.
- Csorgo, M. and P. Revesz (1978). Strong approximations of the quantile process. *The Annals of Statistics* 6(4), 882–894.
- Cuzick, J. (1992a). Efficient estimates in semiparametric additive regression models with unknown error distribution. *The Annals of Statistics* 20(2), 1129–1136.
- Cuzick, J. (1992b). Semiparametric additive regression. *Journal of the Royal Statistical Society. Series B (Methodological)* 45(3), 831–843.

[22] *Athey et al.*

- Doksum, K. (1974). Empirical probability plots and statistical inference for nonlinear models in the two-sample case. *The annals of statistics* 2(2), 267–277.
- Doksum, K. A. and G. L. Sievers (1976). Plotting with confidence: Graphical comparisons of two populations. *Biometrika* 63(3), 421–434.
- Fisher, R. A. (1937). *The design of experiments*. Oliver And Boyd; Edinburgh; London.
- Govindarajulu, Z., L. Le Cam, and M. Raghavachari (1967). Generalizations of theorems of chernoff and savage on the asymptotic normality of test statistics. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pp. 609–638. University of California Press.
- Gupta, S., R. Kohavi, D. Tang, Y. Xu, R. Andersen, E. Bakshy, N. Cardin, S. Chandran, N. Chen, D. Coey, et al. (2019). Top challenges from the first practical online controlled experiments summit. *ACM SIGKDD Explorations Newsletter* 21(1), 20–35.
- Hampel, F. R., E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel (2011). *Robust statistics: the approach based on influence functions*, Volume 196. John Wiley & Sons.
- Hodges Jr, J. L. and E. L. Lehmann (1963). Estimates of location based on rank tests. *The Annals of Mathematical Statistics* 34(2), 598–611.
- Huber, P. J. (1964). Robust estimation of a location parameter. *The Annals of Mathematical Statistics* 35(1), 73–101.
- Huber, P. J. (2011). *Robust statistics*. Springer.
- Imbens, G. W. and D. B. Rubin (2015). *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press.
- Jaekel, L. A. (1971a). Robust estimates of location: Symmetry and asymmetric contamination. *The Annals of Mathematical Statistics* 42(3), 1020–1034.
- Jaekel, L. A. (1971b). Some flexible estimates of location. *The Annals of Mathematical Statistics* 45(2), 1540–1552.
- Klaassen, C. A. (1987). Consistent estimation of the influence function of locally asymptotically linear estimators. *The Annals of Statistics* 15(4), 1548–1562.
- Koenecke, A., M. Powell, R. Xiong, Z. Shen, N. Fischer, S. Huq, A. M. Khalafallah, M. Trevisan, P. Sparen, J. J. Carrero, A. Nishimura, B. Caffo, E. A. Stuart, R. Bai, V. Staedtke, D. L. Thomas, N. Papadopoulos, K. W. Kinzler, B. Vogelstein, S. Zhou, C. Bettgowda, M. F. Konig, B. D. Mensh, J. T. Vogelstein, and S. Athey (2021). Alpha-1 adrenergic receptor antagonists to prevent hyperinflammation and death from lower respiratory tract infection. *eLife* 10, e61700.

- Kohavi, R., D. Tang, and Y. Xu (2020). *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing*. Cambridge University Press.
- Le Cam, L. and G. L. Yang (1988). On the preservation of local asymptotic normality under information loss. *The Annals of Statistics* 16(2), 483–520.
- Lehmann, E. L. and H. J. D’Abrera (1975). *Nonparametrics: statistical methods based on ranks*. Holden-day.
- Lewis, R. A. and J. M. Rao (2015). The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics* 130(4), 1941–1973.
- Linden, L. and J. E. Rockoff (2008). Estimates of the impact of crime risk on property values from megan’s laws. *American Economic Review* 98(3), 1103–27.
- Linden, L. and J. E. Rockoff (2019). Replication data for: Estimates of the impact of crime risk on property values from megan’s laws. *Inter-university Consortium for Political and Social Research (ICPSR)*.
- Neyman, J. (1923/1990). On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science* 5(4), 465–472.
- Rosenbaum, P. R. (1993). Hodges-lehmann point estimates of treatment effect in observational studies. *Journal of the American Statistical Association* 88(424), 1250–1253.
- Rosenbaum, P. R. (2002). *Observational Studies*. Springer.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology* 66(5), 688.
- Stigler, S. M. (1974). Linear functions of order statistics with smooth weight functions. *The Annals of Statistics* 2(4), 676–693.
- Stone, C. J. (1975). Adaptive maximum likelihood estimators of a location parameter. *The Annals of Statistics* 3(2), 267–284.
- Wu, C. J. and M. S. Hamada (2011). *Experiments: planning, analysis, and optimization*, Volume 552. John Wiley & Sons.
- Yu, C. and W. Yao (2017). Robust linear regression: A review and comparison. *Communications in Statistics-Simulation and Computation* 46(8), 6261–6282.

We first provide a general theorem on linear combination of order statistics, which simplifies the results of Chernoff et al. (1967), see also Bickel (1967), Govindarajulu et al. (1967) and Stigler (1974), and then provide the technical details for proving the main results.

6. A general theorem on linear combination of order statistics

Let F be a distribution function, twice differentiable, and $f = F'$. Let $W : [0, 1] \rightarrow \mathcal{R}$ be a function s.t. $W(0) = 0$ and $W(1) = 1$. The C-R condition was first proposed by Csorgo and Revesz (1978).

(C-R) There exists $C_0 > 0$ and $\epsilon \in (0, \frac{1}{2})$ such that

- (a) $F(x)(1 - F(x)) \frac{|f'(x)|}{f^2(x)} \leq C_0$ if $F(x) \in (0, \epsilon)$ or $F(x) \in (1 - \epsilon, 1)$, and C_0 can be replaced by another constant if $F(x) \in [\epsilon, 1 - \epsilon]$.
- (b) $f'(x) \geq 0$ for $x < F^{-1}(\epsilon)$ and $f'(x) \leq 0$ for $x > F^{-1}(1 - \epsilon)$.

(W) There exists $\epsilon_n = o(1)$ such that

$$\int_0^{\epsilon_n} (1 + t^{1-C_0}) |dW(t)| = o(n^{-1/2}) \text{ and } \int_{1-\epsilon_n}^1 (1 + (1-t)^{1-C_0}) |dW(t)| = o(n^{-1/2}) \quad (\text{A.1})$$

$$\int_{\epsilon_n}^{1-\epsilon_n} \frac{1}{f(F^{-1}(t))} dW(t) = O\left(\frac{n^{1/4}}{\log n}\right). \quad (\text{A.2})$$

Theorem A.1. *Let $X_1, \dots, X_n$ be i.i.d. from F , and let $\hat{F}$ and $\hat{F}^{-1}$ be the empirical distribution function and empirical quantile function respectively. Assume that the C-R condition and the above W condition hold and let*

$$\psi(x) = - \int_0^1 \frac{1(F(x) \leq t) - t}{f(F^{-1}(t))} dW(t). \quad (\text{A.3})$$

Then

$$\sqrt{n} \int_0^1 (\hat{F}^{-1}(t) - F^{-1}(t)) dW(t) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \psi(X_i) + o_P(1)$$

which converges in distribution to $\mathcal{N}(0, \sigma_f^2)$ if further $E\psi^2(X_1) < \infty$, where

$$\sigma_f^2 = \int_0^1 \int_0^1 \frac{\min(s, t) - st}{f(F^{-1}(s))f(F^{-1}(t))} dW(s)dW(t). \quad (\text{A.4})$$

When ψ' exists, we may also represent W as below if it is well defined:

$$W(t) = \frac{\int_0^t \psi'(F^{-1}(u)) du}{\int_0^1 \psi'(F^{-1}(u)) du}. \quad (\text{A.5})$$

Remark 1. *Justification of the conditions for a few common scenarios.*

- (a) Gaussian: $f(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$, then $\epsilon = 0.5$ and $C_0 = 1$ meet the C-R condition.

- (b) Cauchy: $f(x) = \frac{1}{\pi(1+x^2)}$, then $\epsilon = 0.5$ and $C_0 = 2$ meet the C-R condition.
- (c) Equal-weight: $W'(t) = 1$, then (A.1) implies $\epsilon_n = o(n^{-1/2})$ for $C_0 \leq 1$ and $\epsilon_n = o(n^{-\frac{1}{2(2-C_0)}})$ for $C_0 \in (1, 2)$. So the equal weight applies to Gaussian. This does not apply to Cauchy, but note that its sample mean does not converge.
- (d) Median: W puts all mass at $t = \frac{1}{2}$. The W condition holds for any f .

6.1. A lemma on the C-R condition

Lemma A.1. *If the C-R condition holds, then for any $t \in (0, \epsilon)$*

$$\frac{1}{f(F^{-1}(t))} \leq K_0 t^{-C_0} \quad (\text{A.6})$$

$$|F^{-1}(t)| \leq K_1 t^{1-C_0} + K_2 \quad (\text{A.7})$$

where K_0, K_1 and K_2 are positive constants, only dependent on C_0 and ϵ . Similar results hold for $t \in (1 - \epsilon, 1)$, omitted for conciseness.

Proof. By the C-R condition, $f'(x) \geq 0$ for any $x < F^{-1}(\epsilon)$, and for any $t \in (0, \epsilon)$,

$$\int_t^\epsilon \frac{f'}{f^2}(F^{-1}(u)) du \leq \int_t^\epsilon \frac{C_0}{u(1-u)} du = C_0 \log \frac{\epsilon}{1-\epsilon} - C_0 \log \frac{t}{1-t} \leq C - C_0 \log t$$

where $C = C_0 \log \frac{\epsilon}{1-\epsilon}$. Since

$$\frac{f'}{f^2}(F^{-1}(u)) = \frac{\partial}{\partial u} \log f(F^{-1}(u)),$$

then

$$\log \frac{f(F^{-1}(\epsilon))}{f(F^{-1}(t))} \leq C - C_0 \log t,$$

which implies (A.6) with $K_0 = \frac{e^C}{f(F^{-1}(\epsilon))}$. Next, (A.7) follows from the inequality below:

$$F^{-1}(\epsilon) - F^{-1}(t) = \int_t^\epsilon \frac{1}{f(F^{-1}(u))} du \leq \int_t^\epsilon K_0 u^{-C_0} du = \frac{K_0}{1-C_0} (\epsilon^{1-C_0} - t^{1-C_0}).$$

□

6.2. Proof of Theorem A.1

Proof. First of all,

$$\int_{\epsilon_n}^{1-\epsilon_n} \left| (\hat{F}^{-1}(t) - F^{-1}(t)) + \frac{\hat{F}(F^{-1}(t)) - t}{f(F^{-1}(t))} \right| \cdot |dW(t)| \leq \Delta_n \cdot \int_{\epsilon_n}^{1-\epsilon_n} \frac{|dW(t)|}{f(F^{-1}(t))}$$

where

$$\Delta_n = \sup_{0 < t < 1} \left| f(F^{-1}(t)) (\hat{F}^{-1}(t) - F^{-1}(t)) + (\hat{F}(F^{-1}(t)) - t) \right|.$$

Under the C-R condition, by Theorem 4 of Csorgo and Revesz (1978), we have

$$\Delta_n = o_P(n^{-3/4} \log n).$$

[26] *Athey et al.*

Under the condition (A.2),

$$\int_{\epsilon_n}^{1-\epsilon_n} \frac{|dW(t)|}{f(F^{-1}(t))} = O_P\left(\frac{n^{1/4}}{\log n}\right).$$

Thus

$$\int_{\epsilon_n}^{1-\epsilon_n} (\hat{F}^{-1}(t) - F^{-1}(t))dW(t) = - \int_{\epsilon_n}^{1-\epsilon_n} \frac{\hat{F}(F^{-1}(t)) - t}{f(F^{-1}(t))} dW(t) + o_P(n^{-1/2}).$$

Next, since

$$\begin{aligned} E \left| \int_0^{\epsilon_n} \frac{\hat{F}(F^{-1}(t)) - t}{f(F^{-1}(t))} dW(t) \right| &\leq 2 \int_0^{\epsilon_n} \frac{t}{f(F^{-1}(t))} |dW(t)| \\ &\leq 2K_0 \int_0^{\epsilon_n} t^{1-C_0} |dW(t)| \\ &= o(n^{-1/2}) \end{aligned}$$

where the last inequality is due to Lemma A.1 and the last equality is due to (A.1), then

$$\int_0^{\epsilon_n} \frac{\hat{F}(F^{-1}(t)) - t}{f(F^{-1}(t))} dW(t) = o_P(n^{-1/2}).$$

Furthermore, by Lemma A.1,

$$\left| \int_0^{\epsilon_n} F^{-1}(t) dW(t) \right| \leq K_1 \int_0^{\epsilon_n} t^{1-C_0} |dW(t)| + K_2 \int_0^{\epsilon_n} |dW(t)| = o(n^{-1/2})$$

where the last equality is due to (A.1). Similarly, we have

$$\int_{1-\epsilon_n}^1 \frac{\hat{F}(F^{-1}(t)) - t}{f(F^{-1}(t))} dW(t) = o_P(n^{-1/2})$$

and

$$\int_{1-\epsilon_n}^1 F^{-1}(t) dW(t) = o(n^{-1/2}).$$

Hence, to prove that

$$\int_0^1 (\hat{F}^{-1}(t) - F^{-1}(t))dW(t) = - \int_0^1 \frac{\hat{F}(F^{-1}(t)) - t}{f(F^{-1}(t))} dW(t) + o_P(n^{-1/2}),$$

it remains to show that

$$\left(\int_0^{\epsilon_n} + \int_{1-\epsilon_n}^1 \right) \hat{F}^{-1}(t) dW(t) = o_P(n^{-1/2}). \quad (\text{A.8})$$

Let $X_{(1)} \leq \dots \leq X_{(n)}$ be the order statistics and $U_{(i)} \equiv F(X_{(i)})$, then $\hat{F}^{-1}(t) = F^{-1}(U_{(\lceil nt \rceil)})$. If $C_0 > 1$, then by Lemma A.1,

$$\left| \int_0^{\epsilon_n} \hat{F}^{-1}(t) dW(t) \right| = \left| \int_0^{\epsilon_n} F^{-1}(U_{(\lceil nt \rceil)}) dW(t) \right| \leq (K_1 + K_2) \int_0^{\epsilon_n} (U_{(\lceil nt \rceil)})^{1-C_0} |dW(t)|.$$

Since $U_{(i)}$ follows the Beta($i, n+1-i$) distribution, one can verify that

$$E(U_{(i)}^a) = \frac{\Gamma(n+1)\Gamma(i+a)}{\Gamma(n+1+a)\Gamma(i)}.$$

Using the bounds for gamma functions (Batir, 2017), i.e. for $x \geq 1$,

$$\sqrt{2\pi}x^xe^{-x}(x^2 + \frac{x}{3} + 0.0496)^{1/4} < \Gamma(x+1) < \sqrt{2\pi}x^xe^{-x}(x^2 + \frac{x}{3} + 0.056)^{1/4},$$

one can get for $i \geq 2-a$,

$$C_1 \cdot (i/n)^a \leq \mathbb{E}U_{(i)}^a \leq C_2 \cdot (i/n)^a$$

where $C_1, C_2 > 0$ are some constants. For $1 \leq i < 2-a$, $\mathbb{E}U_{(i)}^a = \frac{\Gamma(i+a)}{\Gamma(i)}n^{-a}(1+o(1))$. These imply that $\mathbb{E}U_{(\lceil nt \rceil)}^{1-C_0} \leq C_3 \cdot t^{1-C_0}$ for C_3 large enough, and hence

$$\begin{aligned} E\left(\left|\int_0^{\epsilon_n} \hat{F}^{-1}(t)dW(t)\right|\right) &\leq (K_1 + K_2) \int_0^{\epsilon_n} \mathbb{E}U_{(\lceil nt \rceil)}^{1-C_0} |dW(t)| \\ &\leq C_3(K_1 + K_2) \int_0^{\epsilon_n} t^{1-C_0} |dW(t)|. \end{aligned}$$

If $0 < C_0 \leq 1$, then

$$\mathbb{E}\left(\left|\int_0^{\epsilon_n} \hat{F}^{-1}(t)dW(t)\right|\right) \leq (K_1 + K_2) \int_0^{\epsilon_n} |dW(t)|.$$

Therefore, due to (A.1), $\mathbb{E}\left|\int_0^{\epsilon_n} \hat{F}^{-1}(t)dW(t)\right| = o(n^{-1/2})$ for both $C_0 > 1$ and $C_0 \in (0, 1]$, which implies $\int_0^{\epsilon_n} \hat{F}^{-1}(t)dW(t) = o_P(n^{-1/2})$. Similarly, it can be shown that $\int_{1-\epsilon_n}^1 \hat{F}^{-1}(t)dW(t) = o_P(n^{-1/2})$.

Thus (A.8) holds, which concludes the proof. $\square$

7. Proof of Theorem 1

Our proof is based on the sample splitting argument, see (Klaassen, 1987).

7.1. General conditions for the density estimation

Let $X_1, \dots, X_n$ be i.i.d. F with $f \equiv F'$ and $\hat{F}$ be the empirical distribution. Let $\hat{f}$ be an estimate of f from data independent of $\{X_1, \dots, X_n\}$ such that $I(\hat{f}) = I(f) + o_P(1)$ and the following conditions hold for any $\delta_n = O_P(n^{-1/2})$:

$$\int \left(\frac{\hat{f}'}{\hat{f}} - \frac{f'}{f} \right)^2 (x) dF(x) = o_P(1) \tag{A.9}$$

$$\frac{1}{n} \sum_{i=1}^n \left(\frac{\hat{f}'}{\hat{f}} \right)' (X_i + \delta_n) = \frac{1}{n} \sum_{i=1}^n \left(\frac{f'}{f} \right)' (X_i) + o_P(1) = -I(f) + o_P(1). \tag{A.10}$$

Remark 2. These conditions are satisfied by the kernel density estimator under mild conditions, see (Stone, 1975) and (Bickel, 1982).

[28] *Athey et al.*

Let

$$w_f(u) = \frac{1}{I(f)} \left(-\frac{f'}{f} \right)' (F^{-1}(u))$$

where the dependency on F is suppressed for notational simplicity.

Let $\hat{f}$ and $\hat{F}$ be the estimates of f and F respectively, and let

$$w_{\hat{f}}(u) = \left(\frac{\hat{f}'}{\hat{f}} \right)' (\hat{F}^{-1}(u)) \Big/ \int_0^1 \left(\frac{\hat{f}'}{\hat{f}} \right)' (\hat{F}^{-1}(u)) du$$

be the estimate of w_f . Then $\int_0^1 w_{\hat{f}}(u) du = 1$.

We assume that the following conditions hold for $w_{\hat{f}}$:

$$\int_0^1 \frac{|w_{\hat{f}}(u)|}{f(F^{-1}(u))} du = O_P\left(\frac{n^{1/4}}{\log n}\right), \quad (\text{A.11})$$

$$\int_0^1 \int_0^1 \frac{\min(u, v) - uv}{f(F^{-1}(u))f(F^{-1}(v))} \cdot E\left((w_{\hat{f}}(u) - w_f(u))^2\right) dudv = o(1). \quad (\text{A.12})$$

Remark 3. Conditions (A.11) can be achieved by kernel density estimation with proper tail truncation in the spirit of Lemma 6.1 in Bickel (1982).

Remark 4. If $X \sim F$ and $\text{var}(X) < \infty$, then $\int_0^1 \int_0^1 \frac{\min(u, v) - uv}{f(F^{-1}(u))f(F^{-1}(v))} = \text{var}(X)$. See Bickel (1967). This suggests that the requirement (A.12) is not strong.

Lemma A.2. Let $X_1, \dots, X_n$ be i.i.d. F with $f \equiv F'$ and $\hat{F}$ the empirical distribution. Let $w_{\hat{f}}$ be an estimate of w_f independent of $\{X_1, \dots, X_n\}$ and satisfy the above conditions (A.11) and (A.12). Then under the C-R condition,

$$\int_0^1 (\hat{F}^{-1}(u) - F^{-1}(u)) \cdot w_{\hat{f}}(u) du = -\frac{1}{nI(f)} \sum_{i=1}^n \frac{f'}{f}(X_i) + o_P(n^{-1/2}).$$

Proof. Note that

$$\begin{aligned} & \left| \int_0^1 \left(\hat{F}^{-1}(u) - F^{-1}(u) + \frac{\hat{F}(F^{-1}(u)) - u}{f(F^{-1}(u))} \right) \cdot w_{\hat{f}}(u) du \right| \\ & \leq \sup_{0 < u < 1} |f(F^{-1}(u)) \cdot (\hat{F}^{-1}(u) - F^{-1}(u)) + (\hat{F}(F^{-1}(u)) - u)| \cdot \int_0^1 \frac{|w_{\hat{f}}(u)|}{f(F^{-1}(u))} du \\ & = o_P(n^{-3/4} \log n) \cdot \int_0^1 \frac{|w_{\hat{f}}(u)|}{f(F^{-1}(u))} du \end{aligned}$$

where the last equality is due to Theorem 4 of (Csorgo and Revesz, 1978). Thus under the further condition (A.11), we have

$$\int_0^1 \left(\hat{F}^{-1}(u) - F^{-1}(u) + \frac{\hat{F}(F^{-1}(u)) - u}{f(F^{-1}(u))} \right) \cdot w_{\hat{f}}(u) du = o_P(n^{-1/2}).$$

Note that

$$\int_0^1 \frac{\hat{F}(F^{-1}(u)) - u}{f(F^{-1}(u))} w_f(u) du = -\frac{1}{nI(f)} \sum_{i=1}^n \frac{f'}{f}(X_i).$$

Hence it is sufficient to prove that

$$\int_0^1 \frac{\hat{F}(F^{-1}(u)) - u}{f(F^{-1}(u))} \cdot (w_{\hat{f}}(u) - w_f(u)) du = o_P(n^{-1/2}). \quad (\text{A.13})$$

Recall that $w_{\hat{f}}$ is independent of $\hat{F}$, so

$$\begin{aligned} & E \left(\int_0^1 \frac{\hat{F}(F^{-1}(u)) - u}{f(F^{-1}(u))} \cdot (w_{\hat{f}}(u) - w_f(u)) du \right)^2 \\ &= E \int_0^1 \int_0^1 \frac{\hat{F}(F^{-1}(u)) - u}{f(F^{-1}(u))} \cdot \frac{\hat{F}(F^{-1}(v)) - v}{f(F^{-1}(v))} \cdot (w_{\hat{f}} - w_f)(u) \cdot (w_{\hat{f}} - w_f)(v) dudv \\ &= \frac{1}{n} \int_0^1 \int_0^1 \frac{\min(u, v) - uv}{f(F^{-1}(u))f(F^{-1}(v))} \cdot E \left((w_{\hat{f}}(u) - w_f(u))(w_{\hat{f}}(v) - w_f(v)) \right) dudv \\ &\leq \frac{1}{n} \int_0^1 \int_0^1 \frac{\min(u, v) - uv}{f(F^{-1}(u))f(F^{-1}(v))} \cdot E \left((w_{\hat{f}}(u) - w_f(u))^2 \right) dudv \end{aligned}$$

which is $o(\frac{1}{n})$ due to (A.12). Therefore (A.13) holds, which concludes the proof. $\square$

7.2. Proof of Claim 1

The median difference, i.e. $\hat{F}_1^{-1}(\frac{1}{2}) - \hat{F}_0^{-1}(\frac{1}{2})$, is a $\sqrt{n}$ consistent estimate of τ , where $\hat{F}_1^{-1}$ and $\hat{F}_0^{-1}$ are the empirical quantile functions for Y conditional on $Z = 1$ and $Z = 0$, respectively. This follows directly from Theorem A.1.

7.3. Proof of Claim 2

Let $\mathcal{D}_1 = \{(Z_i, Y_i) : 1 \leq i \leq \frac{n}{2}\}$ and $\mathcal{D}_2 = \{(Z_i, Y_i) : \frac{n}{2} + 1 \leq i \leq n\}$ be a sample split. For simplicity, assume that $\sum_{i=1}^{n/2} Z_i = \frac{1}{2}np$ and $\sum_{i=\frac{n}{2}+1}^n Z_i = \frac{1}{2}np$. Let $\hat{f}_{0(j)}$ be the estimate of $f_0 \equiv F'_0$ based on $\mathcal{D}_j$, $j \in \{1, 2\}$, which satisfy the conditions (A.9) and (A.10).

By the first order Taylor expansion at τ , there exists $\bar{\tau}$ between $\tilde{\tau}$ and τ such that

$$\begin{aligned} \hat{\tau}^{if} - \tau &\equiv \tilde{\tau} - \tau + \frac{1}{n} \left(\sum_{i=1}^{n/2} \psi_{\hat{f}_{0(2)}}(Z_i, Y_i; \tilde{\tau}) + \sum_{i=\frac{n}{2}+1}^n \psi_{\hat{f}_{0(1)}}(Z_i, Y_i; \tilde{\tau}) \right) \\ &= A + (\tilde{\tau} - \tau)(1 + B) \end{aligned}$$

where

$$\begin{aligned} A &= \frac{1}{n} \left(\sum_{i=1}^{n/2} \psi_{\hat{f}_{0(2)}}(Z_i, Y_i; \tau) + \sum_{i=\frac{n}{2}+1}^n \psi_{\hat{f}_{0(1)}}(Z_i, Y_i; \tau) \right) \\ B &= \frac{1}{n} \cdot \left(\sum_{i=1}^{n/2} \frac{\partial}{\partial \tau} \psi_{\hat{f}_{0(2)}}(Z_i, Y_i; \bar{\tau}) + \sum_{i=\frac{n}{2}+1}^n \frac{\partial}{\partial \tau} \psi_{\hat{f}_{0(1)}}(Z_i, Y_i; \bar{\tau}) \right). \end{aligned}$$

[30] *Athey et al.*

Since $\hat{f}_{0(2)}$ is independent of $\{(Z_i, Y_i) : 1 \leq i \leq \frac{n}{2}\}$, then

$$\begin{aligned} & E \left(\left(\frac{1}{\sqrt{n}} \sum_{i=1}^{n/2} \left(I(\hat{f}_{0(2)}) \cdot \psi_{\hat{f}_{0(2)}}(Z_i, Y_i; \tau) - I(f_0) \cdot \psi_{f_0}(Z_i, Y_i; \tau) \right) \right)^2 \middle| \mathcal{D}_2 \right) \\ &= \frac{1}{2} E \left(\left(\frac{Z_1}{p} \cdot \left(\frac{\hat{f}'_{0(2)}}{\hat{f}_{0(2)}} - \frac{f'_0}{f_0} \right) (Y_1 - \tau) - \frac{1 - Z_1}{1 - p} \cdot \left(\frac{\hat{f}'_{0(2)}}{\hat{f}_{0(2)}} - \frac{f'_0}{f_0} \right) (Y_1) \right)^2 \middle| \mathcal{D}_2 \right) \\ &= \frac{1}{2p(1-p)} \int \left(\frac{\hat{f}'_{0(2)}}{\hat{f}_{0(2)}} - \frac{f'_0}{f_0} \right)^2 (y) dF_0(y) \end{aligned}$$

which is $o_P(1)$ according to (A.9). Thus

$$\frac{I(\hat{f}_{0(2)})}{\sqrt{n}} \cdot \sum_{i=1}^{n/2} \psi_{\hat{f}_{0(2)}}(Z_i, Y_i; \tau) = \frac{I(f_0)}{\sqrt{n}} \cdot \sum_{i=1}^{n/2} \psi_{f_0}(Z_i, Y_i; \tau) + o_P(1).$$

Furthermore, $I(\hat{f}_{0(2)}) = I(f_0) + o_P(1)$, hence

$$\frac{1}{\sqrt{n}} \sum_{i=1}^{n/2} \psi_{\hat{f}_{0(2)}}(Z_i, Y_i; \tau) = \frac{1}{\sqrt{n}} \sum_{i=1}^{n/2} \psi_{f_0}(Z_i, Y_i; \tau) + o_P(1).$$

Similarly, we have

$$\frac{1}{\sqrt{n}} \sum_{i=1+n/2}^n \psi_{\hat{f}_{0(1)}}(Z_i, Y_i; \tau) = \frac{1}{\sqrt{n}} \sum_{i=1+n/2}^n \psi_{f_0}(Z_i, Y_i; \tau) + o_P(1).$$

Therefore,

$$\begin{aligned} \sqrt{n}A &= \frac{1}{\sqrt{n}} \sum_{i=1}^n \psi_{f_0}(Z_i, Y_i; \tau) + o_P(n^{-1/2}) \\ &\stackrel{d}{\Rightarrow} \mathcal{N}\left(0, \frac{1}{p(1-p)I(f_0)}\right). \end{aligned}$$

From (A.10),

$$\begin{aligned} B &= \frac{1}{I(f_0)} \cdot \frac{1}{n} \sum_{i=1}^n \frac{Z_i}{p} \left(\frac{f'_0}{f_0} \right)' (Y_i - \tau) + o_P(1) \\ &= -1 + o_P(1). \end{aligned}$$

By Slutsky's theorem,

$$\begin{aligned} \sqrt{n}(\hat{\tau}^{if} - \tau) &= \sqrt{n}A + \sqrt{n}(\tilde{\tau} - \tau)(1 + B) \\ &\stackrel{d}{\Rightarrow} \mathcal{N}\left(0, \frac{1}{p(1-p)I(f_0)}\right) \end{aligned}$$

since $\sqrt{n}(\tilde{\tau} - \tau) = O_P(1)$ and $1 + B = o_P(1)$. This concludes the proof.

7.4. Proof of Claim 3

Let $\mathcal{D}_j$, $j \in \{1, 2\}$ be the sample split as above. Let $\hat{F}_{1(j)}$ and $\hat{F}_{0(j)}$ be the empirical distributions of $Y|Z = 1$ and $Y|Z = 0$ respectively based on $\mathcal{D}_j$, for $j \in \{1, 2\}$. Let $w_{\hat{f}_j}$ be as defined earlier and assume that $w_{\hat{f}_j}$ satisfy the conditions of Lemma A.2.

Then the L estimate of τ can be written as:

$$\begin{aligned}\hat{\tau}^{waq} &= \frac{1}{2} \int_0^1 (\hat{F}_{1(1)}^{-1}(u) - \hat{F}_{0(1)}^{-1}(u)) \cdot w_{\hat{f}_{0(2)}}(u) du \\ &\quad + \frac{1}{2} \int_0^1 (\hat{F}_{1(2)}^{-1}(u) - \hat{F}_{0(2)}^{-1}(u)) \cdot w_{\hat{f}_{0(1)}}(u) du.\end{aligned}\tag{A.14}$$

According to Lemma A.2,

$$\begin{aligned}\int_0^1 (\hat{F}_{0(1)}^{-1}(u) - F_0^{-1}(u)) \cdot w_{\hat{f}_{0(2)}}(u) du &= -\frac{2}{n(1-p)I(f_0)} \sum_{i=1}^{n/2} \frac{f'_0}{f_0}(Y_i) \cdot (1 - Z_i) + o_P(n^{-1/2}) \\ \int_0^1 (\hat{F}_{1(1)}^{-1}(u) - F_1^{-1}(u)) \cdot w_{\hat{f}_{0(2)}}(u) du &= -\frac{2}{npI(f_0)} \sum_{i=1}^{n/2} \frac{f'_0}{f_0}(Y_i - \tau) \cdot Z_i + o_P(n^{-1/2})\end{aligned}$$

where the second equality uses the fact that $F_1(\cdot) \equiv F_0(\cdot - \tau)$ and $F_1^{-1}(u) \equiv F_0^{-1}(u) + \tau$. Thus,

$$\begin{aligned}\int_0^1 (\hat{F}_{1(1)}^{-1}(u) - \hat{F}_{0(1)}^{-1}(u)) \cdot w_{\hat{f}_{0(2)}}(u) du &= -\frac{2}{nI(f_0)} \sum_{i=1}^{n/2} \left(\frac{f'_0}{f_0}(Y_i - \tau) \cdot \frac{Z_i}{p} - \frac{f'_0}{f_0}(Y_i) \cdot \frac{1 - Z_i}{1 - p} \right) \\ &\quad + \tau + o_P(n^{-1/2}).\end{aligned}$$

A similar approximation holds for the other term, then by Slutsky's theorem,

$$\hat{\tau}^{waq} = \tau - \frac{1}{n} \sum_{i=1}^n \frac{1}{I(f_0)} \cdot \left(\frac{f'_0}{f_0}(Y_i - \tau) \cdot \frac{Z_i}{p} - \frac{f'_0}{f_0}(Y_i) \cdot \frac{1 - Z_i}{1 - p} \right) + o_P(n^{-1/2}).$$

Hence $\sqrt{n}(\hat{\tau}^{waq} - \tau) \xrightarrow{d} \mathcal{N}\left(0, \frac{1}{p(1-p)I(f_0)}\right)$.

SUPPLEMENTARY MATERIALS FOR:

SEMIPARAMETRIC ESTIMATION OF TREATMENT EFFECTS

IN RANDOMIZED EXPERIMENTS

ATHEY, BICKEL, CHEN, IMBENS & POLLMANN

1. Proof of partial adaptation with (α, β) -trimmed mean**1.1. Main results**

Theorem A.2. Suppose that f_0 falls into the extended Huber family as defined in Section 2.4. If α and β are chosen properly, then for the constant additive treatment effect model, $\hat{\tau}_{\alpha, \beta}$ is an efficient estimate of τ .

Proof. According to Theorem 1, the information bound is given by $\frac{1}{p(1-p)}I^{-1}$. According to Theorem A.1, we have

$$\sqrt{n}(\hat{\tau}_{\alpha, \beta} - \tau) \xrightarrow{d} \mathcal{N}\left(0, p^{-1}\sigma_{f_1}^2 + (1-p)^{-1}\sigma_{f_0}^2\right),$$

where $\sigma_{f_j}^2$ is defined by (A.4) for $j \in \{0, 1\}$, i.e.

$$\sigma_{f_j}^2 = \frac{1}{(1-\alpha-\beta)^2} \int_{\alpha}^{\beta} \int_{\alpha}^{\beta} \frac{\min(s, t) - st}{f_j(F_j^{-1}(s))f_j(F_j^{-1}(t))} ds dt. \quad (\text{A.15})$$

Since $F_1^{-1}(t) - F_0^{-1}(t) \equiv \tau$, then $\sigma_{f_0}^2 = \sigma_{f_1}^2$. Therefore it is sufficient to prove that $\sigma_{f_0}^2 = I^{-1}$. For notational simplicity, below we restrict the proof to $\mu = 0$ and $\sigma = 1$, that is f_0 is parameterized by (k_1, k_2) as (2.16) in Section 2.4 and $f_1(x) \equiv f_0(x - \tau)$.

Let $\alpha = F_0(-k_1)$ and $1 - \beta = F_0(k_2)$.

Then $I = -\int (\frac{f_0'}{f_0})(x) dF_0(x) = \int_{-k_1}^{k_2} dF_0(x) = 1 - (\alpha + \beta)$.

First of all, with integral change of variable,

$$\sigma_{f_0}^2 = \frac{1}{(1-\alpha-\beta)^2} \int_{-k_1}^{k_2} \int_{-k_1}^{k_2} (F_0(\min(x, y)) - F_0(x)F_0(y)) dx dy.$$

Furthermore, it is not hard to verify that $F_0(-k_1) = f_0(-k_1)/k_1$ and $1 - F_0(k_2) = f_0(k_2)/k_2$. Using the fact that $f_0'(x) = -xf_0(x)$ for $x \in (-k_1, k_2)$, then

$$\begin{aligned} \int_{-k_1}^{k_2} F_0(x) dx &= xF_0(x) \Big|_{-k_1}^{k_2} - \int_{-k_1}^{k_2} x f_0(x) dx \\ &= k_2 \cdot (1 - \beta) + k_1 \alpha + \int_{-k_1}^{k_2} df_0(x) \\ &= k_2 \cdot (1 - \beta) + k_1 \alpha + f_0(k_2) - f_0(-k_1) \\ &= k_2 \cdot (1 - \beta) + k_1 \alpha + k_2 \beta - k_1 \alpha \\ &= k_2 \end{aligned}$$

[2] *Athey et al.*

and

$$\begin{aligned}
& \int_{-k_1}^{k_2} \int_{-k_1}^{k_2} (F_0(\min(x, y))) dx dy \\
&= 2 \int_{-k_1}^{k_2} \int_{-k_1}^{k_2} F_0(x) \cdot 1(x \leq y) dx dy \\
&= 2 \int_{-k_1}^{k_2} F_0(x) \cdot (k_2 - x) dx \\
&= - \int_{-k_1}^{k_2} F(x) d(x - k_2)^2 \\
&= \int_{-k_1}^{k_2} (x - k_2)^2 f(x) dx + (k_2 + k_1)^2 \alpha \\
&= \int_{-k_1}^{k_2} (k_2^2 f(x) + 2k_2 f'(x) - x f'(x)) dx + (k_2 + k_1)^2 \alpha \\
&= k_2^2 (1 - \alpha - \beta) + 2k_2 (k_2 \beta - k_1 \alpha) + (k_1 + k_2)^2 \alpha - \int_{-k_1}^{k_2} x f'(x) dx \\
&= k_2^2 + \alpha k_1^2 + \beta k_2^2 - \int_{-k_1}^{k_2} x df(x) \\
&= k_2^2 + \int_{-k_1}^{k_2} f(x) dx \\
&= k_2^2 + 1 - (\alpha + \beta).
\end{aligned}$$

Thus

$$\begin{aligned}
\sigma_{f_0}^2 &= \frac{1}{(1 - (\alpha + \beta))^2} ((k_2^2 + 1 - (\alpha + \beta)) - k_2^2) \\
&= I^{-1}.
\end{aligned}$$

which completes the proof. $\square$

Theorem A.3. *Let*

$$(\hat{\alpha}, \hat{\beta}) = \arg \min_{\alpha_0 \leq \alpha, \beta \leq \alpha_1} (p^{-1} \hat{\sigma}_1^2(\alpha, \beta) + (1 - p)^{-1} \hat{\sigma}_0^2(\alpha, \beta))$$

where $p^{-1} \hat{\sigma}_1^2(\alpha, \beta) + (1 - p)^{-1} \hat{\sigma}_0^2(\alpha, \beta)$ is the estimate of the asymptotic variance as derived in (A.15) and $\hat{\sigma}_j^2$ is the estimate of $\sigma_{f_j}^2$ as given in Proposition A.1. If $0 < \alpha_0, \alpha_1 < 1$ and $\alpha_0 + \alpha_1 < 1$, and $\sigma_1^2(\alpha, \beta)$ has a unique solution w.r.t. (α, β) , which falls into $(\alpha_0, \alpha_1) \times (\alpha_0, \alpha_1)$, then $(\hat{\alpha}, \hat{\beta})$ is a consistent estimate of the optimal (α, β) .

Proof. From Proposition A.1, uniform convergence below holds for $j \in \{0, 1\}$:

$$\sup_{\alpha_0 \leq (\alpha, \beta) \leq \alpha_1} |\hat{\sigma}_j^2(\alpha, \beta) - \sigma_j^2(\alpha, \beta)| = o_P(1).$$

The conclusion holds. $\square$

Remark 5. By Proposition A.2, the limiting objective function $p^{-1} \sigma_1^2(\alpha, \beta) + (1 - p)^{-1} \sigma_0^2(\alpha, \beta)$ is strictly convex at the local neighborhood of the optimal (α, β) .

1.2. Supporting results

Lemma A.3. The influence function (A.3), specialized for the (α, β) -trimmed mean, can be simplified to

$$\begin{aligned}\psi(x) &= \frac{1}{1 - \alpha - \beta}(x - \mu(\alpha, \beta)), & \text{if } x \in (F^{-1}(\alpha), F^{-1}(1 - \beta)) \\ &= \frac{1}{1 - \alpha - \beta}(F^{-1}(1 - \beta) - \mu(\alpha, \beta)), & \text{if } x \geq F^{-1}(1 - \beta) \\ &= \frac{1}{1 - \alpha - \beta}(F^{-1}(\alpha) - \mu(\alpha, \beta)), & \text{if } x \leq F^{-1}(\alpha)\end{aligned}$$

where

$$\mu(\alpha, \beta) = \int_{\alpha}^{1-\beta} F^{-1}(u) du + \alpha F^{-1}(\alpha) + \beta F^{-1}(1 - \beta).$$

The asymptotic variance σ_f^2 , denoted as $\sigma^2(\alpha, \beta)$ for the (α, β) -trimmed mean, can be written as

$$\sigma^2(\alpha, \beta) = \int_0^1 \psi^2(F^{-1}(u)) du.$$

Proof. The proof follows from direct calculation and is omitted. $\square$

Proposition A.1. Assume that f is continuous and that $f(x) > 0$ for any $x \in \mathcal{R}$. Let $\hat{\mu}(\alpha, \beta)$ and $\hat{\sigma}^2(\alpha, \beta)$ be the estimate of $\mu(\alpha, \beta)$ and $\sigma^2(\alpha, \beta)$ respectively as defined in Lemma A.3, by replacing F with $\hat{F}$. Assume that $\sigma^2(\alpha, \beta)$ has a unique minimum w.r.t (α, β) which falls into $(\alpha_0, \alpha_1) \times (\alpha_0, \alpha_1)$, where $(\alpha_0, \alpha_1) \in (0, 1) \times (0, 1)$ with $\alpha_0 + \alpha_1 < 1$. Let

$$(\hat{\alpha}, \hat{\beta}) = \arg \min_{0 < \alpha_0 \leq \alpha, \beta \leq \alpha_1 < 1} \hat{\sigma}^2(\alpha, \beta)$$

be the estimate of the optimal (α, β) . Then $(\hat{\alpha}, \hat{\beta})$ is a consistent estimate of the optimal (α, β) .

Proof. Let $\psi(x)$ be the same as in Lemma A.3, and let $\hat{\psi}$ be an estimate of ψ by replacing F with $\hat{F}$. For $u \in (0, 1)$, define

$$e(u) = \psi(F^{-1}(u)) \text{ and } \hat{e}(u) = \hat{\psi}(\hat{F}^{-1}(u)).$$

Then

$$\begin{aligned}|\hat{\sigma}^2(\alpha, \beta) - \sigma^2| &= \left| \int_0^1 \hat{e}^2(u) du - \int_0^1 e^2(u) du \right| \\ &= \left| \int_0^1 (\hat{e} - e)^2(u) du + 2 \int_0^1 e(u)(\hat{e}(u) - e(u)) du \right| \\ &\leq |\hat{e} - e|_{\infty}^2 + 2|e|_{\infty} \cdot |\hat{e} - e|_{\infty},\end{aligned}$$

where

$$\begin{aligned}|e|_{\infty} &= \sup_{\alpha_0 \leq u \leq 1 - \alpha_1} |e(u)| < \infty \\ |\hat{e} - e|_{\infty} &= \sup_{\alpha_0 \leq u \leq 1 - \alpha_1} |\hat{e}(u) - e(u)| \\ &\leq \frac{2}{1 - \alpha_0 - \alpha_1} \sup_{\alpha_0 \leq u \leq 1 - \alpha_1} |\hat{F}^{-1}(u) - F^{-1}(u)| \\ &= O_P(n^{-1/2}).\end{aligned}$$

[4] *Athey et al.*

Therefore the uniform convergence holds:

$$\sup_{0 < \alpha_0 \leq \alpha, \beta \leq \alpha_1 < 1} |\hat{\sigma}^2(\alpha, \beta) - \sigma^2(\alpha, \beta)| = o_P(1).$$

The conclusion follows since the optimal (α, β) is unique. $\square$

Proposition A.2. *Suppose that f is from the extended Huber family as given by (2.16) in Section 2.4. Then $\sigma^2(\alpha, \beta)$ is strictly convex at the local neighborhood of $(\alpha, \beta) = (F(-k_1), 1 - F(k_2))$.*

Proof. Note that

$$\begin{aligned} (1 - \alpha - \beta)^2 \sigma^2(\alpha, \beta) &= \int_{\alpha}^{1-\beta} \int_{\alpha}^{1-\beta} \frac{\min(s, t) - st}{f(F^{-1}(s))f(F^{-1}(t))} ds dt \\ &\equiv RHS \end{aligned}$$

then

$$\begin{aligned} \frac{\partial RHS}{\partial \alpha} &= -2 \int_{\alpha}^{1-\beta} \frac{\alpha(1-s)}{f(F^{-1}(s))f(F^{-1}(\alpha))} ds \\ &= -\frac{2\alpha}{f(F^{-1}(\alpha))} \int_{\alpha}^{1-\beta} \frac{1-s}{f(F^{-1}(s))} ds \end{aligned}$$

and

$$\frac{\partial RHS}{\partial \beta} = -\frac{2\beta}{f(F^{-1}(1-\beta))} \int_{\alpha}^{1-\beta} \frac{s}{f(F^{-1}(s))} ds.$$

Hence

$$\begin{aligned} \frac{\partial^2 RHS}{\partial \alpha^2} &= \frac{2\alpha}{f(F^{-1}(\alpha))} \frac{1-\alpha}{f(F^{-1}(\alpha))} - 2 \left(\frac{1 - \alpha \frac{f'}{f^2}(F^{-1}(\alpha))}{f(F^{-1}(\alpha))} \right) \int_{\alpha}^{1-\beta} \frac{1-s}{f(F^{-1}(s))} ds \\ &= \frac{2\alpha(1-\alpha)}{f^2(F^{-1}(\alpha))} - \frac{2(f - \alpha f'/f)}{f^2}(F^{-1}(\alpha)) \cdot \int_{\alpha}^{1-\beta} \frac{1-s}{f(F^{-1}(s))} ds \\ \frac{\partial^2 RHS}{\partial \beta^2} &= \frac{2\beta(1-\beta)}{f^2(F^{-1}(1-\beta))} - 2 \left(\frac{f + \beta f'/f}{f^2} \right) (F^{-1}(1-\beta)) \cdot \int_{\alpha}^{1-\beta} \frac{s}{f(F^{-1}(s))} ds \\ \frac{\partial^2 RHS}{\partial \alpha \partial \beta} &= \frac{2\alpha\beta}{f(F^{-1}(\alpha))f(F^{-1}(1-\beta))}. \end{aligned}$$

Now to evaluate the values at

$$\begin{aligned} \alpha &= F(-k_1) = f(-k_1)k_1^{-1} \\ \beta &= 1 - F(k_2) = f(k_2)k_2^{-1}, \end{aligned}$$

we have

$$\begin{aligned} \frac{f'}{f}(F^{-1}(\alpha)) &= k_1 \\ \frac{f'}{f}(F^{-1}(1-\beta)) &= -k_2 \\ \alpha k_1 &= f(-k_1) \\ \beta k_2 &= f(k_2). \end{aligned}$$

Thus

$$\int_{\alpha}^{1-\beta} \frac{s}{f(F^{-1}(s))} ds = \int_{-k_1}^{k_2} F(x) dx = xF(x)|_{-k_1}^{k_2} - \int_{-k_1}^{k_2} x f(x) dx = k_2$$

$$\int_{\alpha}^{1-\beta} \frac{1-s}{f(F^{-1}(s))} ds = (k_2 + k_1) - k_2 = k_1$$

and

$$RHS = 1 - \alpha - \beta$$

$$\frac{\partial RHS}{\partial \alpha} = -2$$

$$\frac{\partial RHS}{\partial \beta} = -2$$

and

$$\frac{\partial^2 RHS}{\partial \alpha^2} = \frac{2(1-\alpha)}{k_1^2 \alpha}$$

$$\frac{\partial^2 RHS}{\partial \beta^2} = \frac{2(1-\beta)}{k_2^2 \beta}$$

$$\frac{\partial^2 RHS}{\partial \alpha \partial \beta} = \frac{2}{k_1 k_2}.$$

Hence

$$\frac{\partial \sigma^2(\alpha, \beta)}{\partial \alpha} = \frac{\partial}{\partial \alpha} \frac{RHS}{(1-\alpha-\beta)^2}$$

$$= \frac{-2}{(1-\alpha-\beta)^2} + RHS \cdot \frac{2}{(1-\alpha-\beta)^3} = 0$$

$$\frac{\partial \sigma^2(\alpha, \beta)}{\partial \beta} = \frac{\partial}{\partial \beta} \frac{RHS}{(1-\alpha-\beta)^2} = \frac{-2}{(1-\alpha-\beta)^2} + RHS \cdot \frac{2}{(1-\alpha-\beta)^3} = 0$$

and the Hessian matrix can be calculated as below:

$$2 \cdot (1, 1)^T (1, 1) \cdot \sigma^2(\alpha, \beta) + (1-\alpha-\beta)^2 \frac{\partial^2 \sigma^2(\alpha, \beta)}{\partial(\alpha, \beta) \partial(\alpha, \beta)} = \begin{pmatrix} \frac{2(1-\alpha)}{k_1^2 \alpha} & \frac{2}{k_1 k_2} \\ \frac{2}{k_1 k_2} & \frac{2(1-\beta)}{k_2^2 \beta} \end{pmatrix}$$

so

$$\frac{\partial^2 \sigma^2(\alpha, \beta)}{\partial(\alpha, \beta) \partial(\alpha, \beta)} = \frac{2}{(1-\alpha-\beta)^2} \begin{pmatrix} \frac{(1-\alpha)}{k_1^2 \alpha} - \frac{1}{1-\alpha-\beta} & \frac{1}{k_1 k_2} - \frac{1}{1-\alpha-\beta} \\ \frac{1}{k_1 k_2} - \frac{1}{1-\alpha-\beta} & \frac{(1-\beta)}{k_2^2 \beta} - \frac{1}{1-\alpha-\beta} \end{pmatrix}.$$

Note that

$$\int_{-k_1}^{k_2} x^2 f(x) dx = \int_{-k_1}^{k_2} -x f'(x) dx$$

$$= \int_{-k_1}^{k_2} f(x) dx - x f(x)|_{-k_1}^{k_2}$$

$$= 1 - \alpha - \beta - (k_2^2 \beta + k_1^2 \alpha)$$

[6] *Athey et al.*

that is,

$$\begin{aligned} 1 - \alpha - \beta &= k_1^2 \alpha + k_2^2 \beta + \int_{-k_1}^{k_2} x^2 f(x) dx \\ &= k_1^2 \alpha + k_2^2 \beta + (1 - \alpha - \beta) ES^2 \end{aligned}$$

where S is defined on $x \in (-k_1, k_2)$ with density $\frac{f(x)}{1 - \alpha - \beta}$. Then

$$\begin{aligned} ES &= (1 - \alpha - \beta)^{-1} \int_{-k_1}^{k_2} x f(x) dx \\ &= (1 - \alpha - \beta)^{-1} \int_{-k_1}^{k_2} -f'(x) dx \\ &= \frac{k_1 \alpha - k_2 \beta}{1 - \alpha - \beta} \end{aligned}$$

So the determinant of $\frac{\partial^2 \sigma^2(\alpha, \beta)}{\partial(\alpha, \beta) \partial(\alpha, \beta)} \times \frac{(1 - \alpha - \beta)^2}{2}$ is

$$\begin{aligned} & \left(\frac{(1 - \alpha)}{k_1^2 \alpha} - \frac{1}{1 - \alpha - \beta} \right) \left(\frac{(1 - \beta)}{k_2^2 \beta} - \frac{1}{1 - \alpha - \beta} \right) - \left(\frac{1}{k_1 k_2} - \frac{1}{1 - \alpha - \beta} \right)^2 \\ &= \frac{(1 - \alpha)}{k_1^2 \alpha} \frac{(1 - \beta)}{k_2^2 \beta} - \frac{1}{k_1^2 k_2^2} + \frac{1}{1 - \alpha - \beta} \left(\frac{2}{k_1 k_2} - \frac{(1 - \alpha)}{k_1^2 \alpha} - \frac{(1 - \beta)}{k_2^2 \beta} \right) \\ &= \frac{k_1^2 \alpha + k_2^2 \beta + (1 - \alpha - \beta) ES^2}{k_1^2 k_2^2 \alpha \beta} - \frac{1}{1 - \alpha - \beta} \cdot \frac{k_1^2 \alpha (1 - \beta) + k_2^2 \beta (1 - \alpha) - 2 k_1 k_2 \alpha \beta}{k_1^2 k_2^2 \alpha \beta} \\ &= \frac{(1 - \alpha - \beta)^2 ES^2 - (k_1^2 \alpha^2 + k_2^2 \beta^2 - 2 k_1 k_2 \alpha \beta)}{(1 - \alpha - \beta) k_1^2 k_2^2 \alpha \beta} \\ &= \frac{1 - \alpha - \beta}{k_1^2 k_2^2 \alpha \beta} \cdot \text{var}(S) \end{aligned}$$

which is greater than 0 since $\alpha + \beta < 1$. This completes the proof. $\square$

2. Proof of Theorem 2

2.1. Proof of Lemma 1

Let $g(y, \theta) = \frac{\partial}{\partial \theta} \log f_1(y, \theta)$, where $f_1(y, \theta)$ is the same as in (3.18).

In order to find the projection of $\dot{\ell}$ on the orthocomplement of the $\dot{P}_f$, we need to find Q which minimizes

$$\epsilon(Q) \equiv E \left(\dot{\ell}(Y, Z, \theta) - (ZQ(h^{-1}(Y, \theta), \theta) + (1 - Z)Q(Y, \theta)) \right)^2.$$

Since $Z \in \{0, 1\}$ and $\dot{\ell}(y, z, \theta) = z \cdot g(y, \theta)$, we have

$$\begin{aligned} \epsilon(Q) &= E \left(Z \cdot (g(Y, \theta) - Q(h^{-1}(Y, \theta), \theta))^2 \right) + E \left((1 - Z) \cdot Q^2(Y, \theta) \right) \\ &= p \cdot E \left((g(Y, \theta) - Q(h^{-1}(Y, \theta), \theta))^2 | Z = 1 \right) + (1 - p) \cdot E \left(Q^2(Y, \theta) | Z = 0 \right) \\ &= p \cdot E \left((g(h(Y, \theta), \theta) - Q(Y, \theta))^2 | Z = 0 \right) + (1 - p) \cdot E \left(Q^2(Y, \theta) | Z = 0 \right) \end{aligned}$$

where the last equality is due to $P(Y \leq h(y, \theta) | Z = 1) = P(Y \leq y | Z = 0)$. Therefore,

$$\epsilon(Q) = E \left((Q(Y, \theta) - p \cdot g(h(Y, \theta), \theta))^2 | Z = 0 \right) + \text{const}$$

where $const$ does not depend on Q . This leads to

$$Q(y, \theta) = p \cdot g(h(y, \theta), \theta).$$

Note that this Q satisfies $\int Q(y, \theta) f_0(y) dy = 0$ since

$$\begin{aligned} E \frac{\partial \ell}{\partial \theta}(Y, Z, \theta) &= E(Zg(Y, \theta)) \\ &= p \cdot E(g(Y, \theta) | Z = 1) \\ &= p \cdot E(g(h(Y, \theta), \theta) | Z = 0) \end{aligned}$$

must be 0. Thus

$$\begin{aligned} \dot{\ell}^*(y, z, \theta) &= z \cdot g(y, \theta) - (z \cdot p \cdot g(y, \theta) + (1 - z) \cdot p \cdot g(h(y, \theta), \theta)) \\ &= (1 - p) \cdot z \cdot g(y, \theta) - p \cdot (1 - z) \cdot g(h(y, \theta), \theta). \end{aligned}$$

and

$$\begin{aligned} E(\dot{\ell}^*(Y, Z, \theta))^2 &= (1 - p)^2 E(Z \cdot g^2(Y, \theta)) + p^2 E((1 - Z) \cdot g^2(h(Y, \theta), \theta)) \\ &= p(1 - p) \int g^2(h(y, \theta), \theta) f_0(y) dy. \end{aligned}$$

2.2. Proof of Theorem 2

The proof is similar to the sample-splitting argument for Claim 3 of Theorem 1 and thus is omitted for conciseness.

3. Details on the Empirical Implementation

3.1. Density Estimation

Our estimators rely on estimates of the density and its first and second derivative.

Using data X (for instance half the sample of control observations for the estimator with sample splitting), we estimate $f^{(k)}(x)$ at $x \in \mathbb{X} \equiv \{X_{(\lceil n_d \frac{k}{1000} \rceil)} \mid k = 1, \dots, 999\}$ where n_d is the number of observations in the data X used for estimating densities. Estimating the density only at points in the density-estimation sample avoids estimating densities of 0 in the tails where there are few observations especially when using sample-splitting. We use adaptive kernel estimates (Silverman 1986, ch. 5.3) and a triweight kernel, which allows estimation of the first and second derivative of the density. Density estimation proceeds in two steps.

First, we obtain an initial estimate of the density at each point $x \in \mathbb{X}$ as

$$\tilde{f}(x) = \frac{1}{n_d h} \sum_{i=1}^{n_d} K\left(\frac{x - X_i}{h}\right)$$

where $K(u)$ is the triweight kernel

$$K(u) = \frac{35}{32} (1 - u^2)^3 \cdot 1_{|u| \leq 1}$$

and h is the bandwidth, chosen as

$$h = 3.15 \cdot \hat{\sigma}_X \cdot n_d^{-1/5}$$

[8] *Athey et al.*

where 3.15 is the Silverman rule-of-thumb constant for the triweight kernel and

$$\hat{\sigma}_X = \frac{\text{quantile}(X, 0.95) - \text{quantile}(X, 0.05)}{2 \cdot 1.6449}$$

is an estimate of the standard deviation of X if X is normally distributed but less affected by outliers than the conventional unbiased estimator of the variance.

Second, we obtain our final estimate of the density by calculating an *adaptive* bandwidth. Intuitively, in regions of high density, we can choose a small bandwidth that lowers bias without incurring much of a cost in terms of variance. In regions of low density, however, a larger bandwidth is needed to have sufficiently many data points near x to reign in the variance. Following Silverman (1986, ch. 5.3.1), define the local bandwidth factors λ_i for X_i as

$$\lambda_i = \left(\frac{\tilde{f}(X_i)}{g} \right)^{-\alpha}$$

where g is the geometric mean of the $\tilde{f}(X_i)$:

$$g = \left(\prod_{i=1}^{n_d} \tilde{f}(X_i) \right)^{1/n_d}.$$

Then the estimate of the density is

$$\hat{f}(x) = \frac{1}{n_d} \sum_{i=1}^{n_d} \frac{1}{h\lambda_i} K\left(\frac{x - X_i}{h\lambda_i}\right).$$

To estimate the first and second derivative of the density, we take λ_i as above, and estimate

$$\begin{aligned} \hat{f}'(x) &= \frac{1}{n_d} \sum_{i=1}^{n_d} \frac{1}{(h_1\lambda_i)^2} K'\left(\frac{x - X_i}{h_1\lambda_i}\right) \\ \hat{f}''(x) &= \frac{1}{n_d} \sum_{i=1}^{n_d} \frac{1}{(h_2\lambda_i)^3} K''\left(\frac{x - X_i}{h_2\lambda_i}\right) \end{aligned}$$

with h_1 and h_2 the Silverman rule-of-thumb bandwidth for estimating the derivatives of the density using a triweight kernel

$$\begin{aligned} h_1 &= 2.83 \cdot \hat{\sigma}_X \cdot n_d^{-1/7} \\ h_2 &= 2.70 \cdot \hat{\sigma}_X \cdot n_d^{-1/9} \end{aligned}$$

and derivatives of the triweight kernel K

$$\begin{aligned} K'(u) &= \frac{3 \cdot 35}{32} (1 - u^2)^2 \cdot (-2u) \cdot 1_{|u| \leq 1} \\ K''(u) &= \frac{2 \cdot 3 \cdot 35}{32} (1 - u^2) \cdot (5u^2 - 1) \cdot 1_{|u| \leq 1} \end{aligned}$$

Then our estimates of the derivatives of the *log* density are, for $x \in \mathbb{X}$,

$$\begin{aligned} \frac{\partial \widehat{\ln f(x)}}{\partial x} &= \frac{\hat{f}'(x)}{\hat{f}(x)} \\ \frac{\partial^2 \widehat{\ln f(x)}}{\partial x^2} &= \frac{\hat{f}(x)\hat{f}''(x) - (\hat{f}'(x))^2}{(\hat{f}(x))^2} \end{aligned}$$

When we need to evaluate the second derivative at some other point $x \notin \mathbb{X}$, we use linear interpolation and nearest-neighbor extrapolation of the estimated derivative of the log density at the points in $\mathbb{X}$. In practice, this piecewise-linear approximation causes negligible error because areas where $\mathbb{X}$ is not tightly spaced have low density by construction and typically very little curvature.

3.2. Efficient Influence Function Estimator

The efficient influence function estimator $\hat{\theta}^{if}$ solves

$$\frac{1}{n_1} \sum_{i=1}^{n_1} \widehat{\frac{\partial \ln f}{\partial x}}(Y_i^{(1)} - \hat{\theta}^{if}) - \frac{1}{n_0} \sum_{i=1}^{n_0} \widehat{\frac{\partial \ln f}{\partial x}}(Y_i^{(0)}) = 0 \quad (\text{A.16})$$

where the estimates $\widehat{\frac{\partial \ln f}{\partial x}}$ either come from the full sample or we use sample-splitting as described in Section 2. Standard root-finding algorithms appear to work well for finding the solution $\hat{\theta}^{if}$ to equation A.16 when initialized at reasonable guesses such as the weighted average quantiles estimator (see below) or the difference in medians.

We estimate standard errors as follows: First, calculate

$$\hat{\mathcal{I}} = -\frac{1}{n_0} \sum_{i=1}^{n_0} \widehat{\frac{\partial^2 \ln f}{\partial x^2}}(Y_i^{(0)})$$

where $Y_i^{(0)}$ for $i = 0, \dots, n_0$ are the control observations. Then an estimate of the approximate finite sample variance of $\hat{\theta}^{if}$ is

$$\widehat{\text{var}}(\hat{\theta}^{if}) = \frac{1}{p(1-p)} \frac{1}{\hat{\mathcal{I}}} \frac{1}{n}$$

where as before $p = \frac{n_1}{n_0+n_1}$ is the fraction of treated observations. This estimator of the variance relies on the correct specification of the additive shift model.

3.3. Weighted Average Quantile Estimator

For the weighted average quantiles estimator, we calculate, assuming $n_0 \geq n_1$

$$\hat{\theta}^{waq} = \sum_{i=1}^{n_0} w\left(\frac{i}{n_0+1}\right) (Y_{(\lceil n_1 \frac{i}{n_0+1} \rceil)}^{(1)} - Y_{(i)}^{(0)}) \quad (\text{A.17})$$

where for $u \in \{i/(n_0+1) : i = 1, \dots, n_0\}$

$$\begin{aligned} \tilde{w}(u) &= \widehat{\frac{\partial^2 \ln f}{\partial x^2}}(\text{quantile}(Y^{(0)}, u)) \\ w(u) &= \begin{cases} c \cdot \tilde{w}(u) & \text{if } \left| \frac{\tilde{w}(u) / \sum_{i=1}^{n_0} \tilde{w}(\frac{i}{n_0+1})}{\hat{f}(\hat{F}^{-1}(u)) \cdot \text{MAD}(\hat{f})} \right| < \log(\log(n)) / \log(n) \cdot n^{1/4} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Here c is for normalization such that $\sum_{i=1}^{n_0} w(\frac{i}{n_0+1}) = 1$, and $\text{quantile}(Y^{(0)}, u)$ is the u quantile of all the control observations for the full-sample estimator. $\hat{f}(\hat{F}^{-1}(u))$ is an ad-hoc estimate of the density at the u quantile calculated as $\hat{f}(\hat{F}^{-1}(i/(n_0+1))) = (1/n_0) / \max(\Delta Y_{(i)}^{(0)}, \Delta Y_{(\lceil n_1 \frac{i}{n_0+1} \rceil)}^{(1)})$

[10] *Athey et al.*

with Δ the first difference operator. The factor $\text{MAD}(\hat{f})$, i.e. the median absolute deviation of the data (multiplied by 1.4826 such that it estimates the standard deviation if the data are normally distributed), is used so that the truncation condition is both shift-invariant and scale-invariant. The truncation of the weights $\tilde{w}$ addresses violations of condition (A.11) in the tails by truncating observations with excessive (normalized) weight relative to the density at the quantile. In the simulations, it improves the performance for the Cauchy distribution for some samples with extreme outliers, and otherwise has no effect.

For a sample-splitting estimator, we randomly split the samples of treated and control in half, use the first half to estimate $w(u)$ (the densities and $\text{quantile}(Y^{(0)}, u)$), and then evaluate equation A.17 for the second half of the sample given w . Reverse the roles of the sample splits and average the two estimators, as described in the main text.

Under correct specification, $\hat{\theta}^{waq}$ and $\hat{\theta}^{if}$ are asymptotically equivalent, so the standard errors based on $\hat{\mathcal{I}}$ are appropriate for $\hat{\theta}^{waq}$ as well.