Published in final edited form as:

Stat Methods Med Res. 2019 April ; 28(4): 1094–1104. doi:10.1177/0962280217747130.

Estimating scaled treatment effects with multiple outcomes

Edward H Kennedy¹, Shreya Kangovi², and Nandita Mitra³

¹Department of Statistics, Carnegie Mellon University, Pittsburgh, PA, USA

²Division of General Internal Medicine, University of Pennsylvania, Pittsburgh, PA, USA

³Department of Biostatistics & Epidemiology, University of Pennsylvania, Pittsburgh, PA, USA

Abstract

In classical study designs, the aim is often to learn about the effects of a treatment or intervention on a single outcome; in many modern studies, however, data on multiple outcomes are collected and it is of interest to explore effects on multiple outcomes simultaneously. Such designs can be particularly useful in patient-centered research, where different outcomes might be more or less important to different patients. In this paper, we propose scaled effect measures (via potential outcomes) that translate effects on multiple outcomes to a common scale, using mean-variance and median-interquartile range based standardizations. We present efficient, nonparametric, doubly robust methods for estimating these scaled effects (and weighted average summary measures), and for testing the null hypothesis that treatment affects all outcomes equally. We also discuss methods for exploring how treatment effects depend on covariates (i.e., effect modification). In addition to describing efficiency theory for our estimands and the asymptotic behavior of our estimators, we illustrate the methods in a simulation study and a data analysis. Importantly, and in contrast to much of the literature concerning effects on multiple outcomes, our methods are nonparametric and can be used not only in randomized trials to yield increased efficiency, but also in observational studies with high-dimensional covariates to reduce confounding bias.

Keywords

Causal inference; doubly robust; multivariate outcomes; outcome-wide analysis; policy evaluation

1 Introduction

In classical study designs, the aim is often to learn about the effects of a treatment or intervention on a single outcome; in many modern studies, however, data on multiple outcomes are collected and it is of interest to explore effects on multiple outcomes simultaneously. Such designs are particularly important in patient-centered research, for

Reprints and permissions: sagepub.co.uk/journalsPermissions.nav

Corresponding author: Edward H Kennedy, Carnegie Mellon University, Baker Hall, Pittsburgh, PA 15213-3815, USA. edward@stat.cmu.edu.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Supplement Material

Supplementary material is available for this article online

example, where different outcomes might be more or less important to different patients,^{1,2} and more generally in prioritizing public health recommendations.³

There has been varied and relatively extensive discussion in the literature over the past few decades about estimating treatment effects on multiple outcomes,^{4–10} including estimating scaled effects,^{11,12} which is a major focus of this paper. However, most of the aforementioned work requires strong parametric assumptions and is geared toward randomized trials rather than observational studies, which can require adjustment for high-dimensional confounders. In contrast, we consider nonparametric doubly robust methods for estimation and hypothesis testing of scaled treatment effects on multiple outcomes. In particular, we translate effects to a common scale with mean-variance and median-interquartile range based standardizations, which are constructed within an explicitly causal potential outcomes framework. Our work is a response to recent proposals by VanderWeele³ and others to spend more effort exploring effects of interventions on multiple outcomes simultaneously, rather than using the classical one-outcome-at-a-time approach. Importantly, our work is designed to accommodate modern studies that include complex covariate information, which can be leveraged for efficiency gains or to reduce confounding bias (or both).

The setup of the paper is as follows. In Sections 3.1 and 3.2, we present efficient doubly robust methods for estimating our proposed scaled effects (along with weighted average summary measures in Section 3.5), as well as methods for testing the null hypothesis that treatment affects all outcomes equally in Section 3.3. We also discuss approaches for exploring how treatment effects vary with covariates (i.e., effect modification) in Section 3.4. In addition to describing efficiency theory for our estimands and the asymptotic behavior of our estimators, in Section 4, we illustrate the methods in a simulation study and in Section 5, we apply them to a recently conducted trial evaluating the effect of community health workers on various health outcomes in a low income population.

2 Setup

2.1 Data and notation

Suppose we observe an independent and identically distributed sample $(\mathbf{Z}_1, \dots, \mathbf{Z}_n)$, where each observation $\mathbf{Z} = (\mathbf{X}, A, \mathbf{Y})$ consists of a vector of p covariates $\mathbf{X} = (X_1, \dots, X_p)$, a binary treatment A , and a vector of K outcomes $\mathbf{Y} = (Y_1, \dots, Y_K)$. We characterize treatment effects using potential outcome notation,¹³ letting $\mathbf{Y}^a = (Y_1^a, \dots, Y_K^a)$ denote the outcome vector that would have been observed under treatment level a .

We use $\mathbb{P}$ to denote the distribution of $\mathbf{Z} = (\mathbf{X}, A, \mathbf{Y})$, and write expectations under $\mathbb{P}$ with usual $\mathbb{E}$ operator notation. For a generic random variable U , we define standard deviations as usual with $\text{SD}(U) = \sqrt{\mathbb{E}(U^2) - \mathbb{E}(U)^2}$. We use $\mathbb{P}_n$ to denote the empirical measure so that sample averages can be written as $\frac{1}{n} \sum_i f(\mathbf{Z}_i) = \mathbb{P}_n\{f(\mathbf{Z})\}$. Finally, we use the following notation to simplify the presentation

$$\pi(a|\mathbf{x}) = \mathbb{P}(A = a|\mathbf{X} = \mathbf{x})$$

$$\mu_k(\mathbf{x}, a) = \mathbb{E}(Y_k|\mathbf{X} = \mathbf{x}, A = a)$$

$$\eta_k(\mathbf{x}, a) = \mathbb{E}(Y_k^2|\mathbf{X} = \mathbf{x}, A = a).$$

2.2 Identification

Throughout this paper, we consider estimating quantities defined in terms of the distributions of potential outcomes $\mathbf{Y}^a$ for $a = 0, 1$. Since potential outcomes are not observed directly, we need identifying assumptions to express estimands of interest in terms of the estimable observed data distribution. We consider the usual ignorability or “no unmeasured confounding” setting, in which the following assumptions hold for $a = 0, 1$:

Assumption 1. Consistency: $A = a$ implies $\mathbf{Y} = \mathbf{Y}^a$.

Assumption 2. Positivity: $\mathbb{P}\{0 < \pi(a|\mathbf{X}) < 1\} = 1$ for all a .

Assumption 3. Exchangeability: $A \perp\!\!\!\perp \mathbf{Y}^a|\mathbf{X}$.

Assumptions 1–3 can hold by design in a randomized trial, since treatment A is under the control of investigators. However, in observational studies these assumptions can be violated, and are generally untestable (apart from positivity). Consistency means potential outcomes are defined uniquely by subjects’ own treatment levels (this can be violated in, e.g., vaccine studies). Positivity means treatment is not assigned deterministically for any subjects, regardless of covariates. Exchangeability means treatment is as good as randomized (within covariate strata) since it is unrelated to potential outcomes once we condition on covariates. Exchangeability requires either external randomization of treatment, or else the collection of sufficiently many relevant covariates.

Assumptions 1–3 have been discussed at length elsewhere, and it is well-known that they imply

$$\mathbb{P}(\mathbf{Y}^a \leq \mathbf{y}|\mathbf{X}) = \mathbb{P}(\mathbf{Y} \leq \mathbf{y}|\mathbf{X}, A = a),$$

that is, the conditional distribution of potential outcomes under $A = a$ (given covariates) equals the conditional distribution of observed outcomes (given covariates) among those for whom $A = a$ observationally. For example, this fact also implies that $\mathbb{E}(\mathbf{Y}^a) = \mathbb{E}\{\mathbb{E}(\mathbf{Y}|\mathbf{X}, A = a)\}$ and similarly for other marginal quantities.

3 Methodology

In this section, we present scaled treatment effect parameters and estimators, discuss how to test for differential effects across multiple outcomes, give extensions for exploring how

treatment effects vary with covariates, and finally present weighted average measures that can provide a scalar summary of multivariate effects.

3.1 Scaled average effects

We start by presenting a scaled treatment effect parameter (using mean-variance standardization), discuss corresponding semiparametric efficiency theory, give doubly robust and locally efficient estimators, and describe asymptotic properties.

As noted for example by Lin et al.¹¹ and Roy et al.,¹² usual outcome-specific effects cannot be compared directly in studies with multiple outcomes measured on different scales. For example, average differences of the form $\mathbb{E}(Y_k^1 - Y_k^0)$ will generally have different and noncomparable units (e.g., kilograms for $k = 1$ and millimeters of mercury for $k = 2$). Thus, to generate a unitless measure of effect, we propose a simple standardization by the standard deviation of outcomes Y_k^0 under control. Note that this is most useful in settings where $A = 0$ represents a meaningful control group (e.g., standard of care), rather than an alternative and potentially comparable treatment option.

Specifically, we characterize effects on outcome k with the scaled effect measure

$$\psi_k = \frac{\mathbb{E}(Y_k^1 - Y_k^0)}{sd(Y_k^0)}. \quad (1)$$

This effect measure captures the mean difference in outcomes under treatment versus control, expressed in terms of the standard deviation under control. Thus, $\psi_k = 1$ indicates that treatment increases outcomes by one standard deviation, on average, of what they would have been under control; similarly, $\psi_k = 2$ means treatment increases outcomes by two standard deviations on average, and $\psi_k = -1$ means treatment decreases outcomes by one standard deviation. Scaled effect measures have played an important role in studies with multiple outcomes,^{11,12} however, so far they have only been proposed within the context of parametric models. Our work can thus be viewed as a nonparametric extension, which also admits doubly robust estimators.

We can use results from semiparametric theory^{14–16} to construct optimal estimators for ψ_k under minimal assumptions about the distribution of the data $\mathbb{P}$. We refer to Kennedy¹⁷ for a review. First define, for $a = 0, 1$

$$\phi_{ak}(\mathbf{Z}; \pi, \mu) = \frac{\mathbb{1}(A = a)}{\pi(a|\mathbf{X})} \{Y_k - \mu_k(\mathbf{X}, a)\} + \mu_k(\mathbf{X}, a) \quad (2)$$

$$\phi_{2k}(\mathbf{Z}; \pi, \eta) = \frac{\mathbb{1}(A = 0)}{\pi(0|\mathbf{X})} \{Y_k^2 - \eta_k(\mathbf{X}, 0)\} + \eta_k(\mathbf{X}, 0) \quad (3)$$

as components of the efficient influence functions for $\mathbb{E}(Y_k^a)$ and $\mathbb{E}\{(Y_k^0)^2\}$, respectively. Then, given estimators $(\hat{\pi}, \hat{\mu}, \hat{\eta})$ of the nuisance functions, the estimator

$$\hat{\psi}_k = \frac{\mathbb{P}_n\{\phi_{1k}(\mathbf{Z}; \hat{\pi}, \hat{\mu}) - \phi_{0k}(\mathbf{Z}; \hat{\pi}, \hat{\mu})\}}{\sqrt{\mathbb{P}_n\{\phi_{2k}(\mathbf{Z}; \hat{\pi}, \hat{\mu})\} - [\mathbb{P}_n\{\phi_{0k}(\mathbf{Z}; \hat{\pi}, \hat{\mu})\}]^2}} \quad (4)$$

is doubly robust and locally efficient, under nonparametric models as well as models that put some (e.g., parametric) restrictions on the treatment mechanism π . We will now discuss these properties in more detail, proofs of which are given in the online Supplementary Materials.

Double robustness is a very important property that has been discussed in detail before.^{18,19} One important consequence of double robustness is that analysts have two chances at obtaining a consistent estimator. For example, in our case, the estimator $\hat{\psi}_k$ is consistent for its target ψ_k as long as either of the nuisance estimators $\hat{\pi}$ or $(\hat{\mu}, \hat{\eta})$ are consistent, even if one of $\hat{\pi}$ or $(\hat{\mu}, \hat{\eta})$ is misspecified. In particular, this means consistency of $\hat{\psi}_k$ is guaranteed in a randomized trial, since there π is known and thus can be estimated consistently under no assumptions. We prove that our estimator is doubly robust in Section 1 of the Supplementary Materials.

Another crucially important property of doubly robust estimators is that they can attain fast parametric $\sqrt{n}$ rates of convergence even after machine learning-based covariate adjustment.²⁰ This is not the case for most standard plug-in estimators, which typically inherit slower than $\sqrt{n}$ convergence rates from their nuisance estimators.²¹ In contrast, our proposed estimator $\hat{\psi}_k$ will be $\sqrt{n}$ -consistent and asymptotically normal even if the nuisance functions $(\hat{\pi}, \hat{\mu}, \hat{\eta})$ are estimated flexibly with nonparametric or machine learning methods, as long as these nuisance estimators converge at faster than $n^{1/4}$ rates (and under some empirical process conditions, which can be avoided with sample splitting).

In particular, under the above conditions and other standard regularity conditions given in the online Supplementary Materials, $\hat{\psi}_k$ is asymptotically normal

$$\sqrt{n}(\hat{\psi}_k - \psi_k) \rightsquigarrow N(0, \sigma_k^2), \quad (5)$$

with asymptotic variance σ_k^2 equal to the variance of the efficient influence function, which is given by $\phi_k(\mathbf{Z}; \pi, \mu, \eta)$ defined as

$$\frac{\phi_{1k}(\mathbf{Z}; \pi, \mu) - \phi_{0k}(\mathbf{Z}; \pi, \mu)}{sd(Y_k^0)} - \psi_k \left[\frac{\phi_{2k}(\mathbf{Z}; \pi, \eta) + \mathbb{E}\{(Y_k^0)^2\} - 2\mathbb{E}(Y_k^0)\phi_{0k}(\mathbf{Z}; \pi, \mu)}{2sd(Y_k^0)^2} \right] \quad (6)$$

Thus, Wald-type confidence intervals for ψ_k can be constructed by estimating the asymptotic variance σ_k^2 with the empirical variance of the estimated efficient influence function values (obtained by replacing unknown quantities in equation (6) with estimates). The asymptotic normality result (equation (5)) is proved in Section 2 of the online Supplementary Materials, using empirical process theory^{22,23} to allow for flexible nonparametric estimation of $(\hat{\pi}, \hat{\mu}, \hat{\eta})$. We prove that equation (6) is in fact the efficient influence function in Section 3 of the online Supplementary Materials, which (given the asymptotic normality result) implies that $\hat{\psi}_k$ is locally efficient. In particular, $\hat{\psi}_k$ is locally semiparametric efficient under models that put at most some restrictions on the treatment mechanism (including, e.g., nonparametric models, as well as models in which the treatment mechanism is known).

Finally, in Section 6 of the online Supplementary Materials we give easily implementable R code for computing $\hat{\psi}_k$ based on estimating equations and targeted maximum likelihood (TMLE).^{20,24} TMLE uses specially constructed nuisance estimates $(\hat{\mu}^*, \hat{\eta}^*)$ that insure that resulting estimators of the numerator and denominator of equation (4) respect the bounds of the parameter space (e.g., the numerator must lie between $[-1, 1]$ when $Y_k \in [0, 1]$), which can improve finite-sample properties and lessen the impact of extreme propensity scores. The provided code also calculates confidence intervals based on the approach described in the previous paragraph.

3.2 Scaled quantile effects

In some cases, the mean and variance are not useful measures of centrality and spread (e.g., for distributions that are highly skewed), in which case the standardization given in equation (1) may not be most appropriate. An alternative quantile-based standardization that is immune to such concerns is given by

$$\psi_k^q = \frac{\mathbb{M}(Y_k^1) - \mathbb{M}(Y_k^0)}{iqr(Y_k^0)}, \quad (7)$$

where for arbitrary random variable U with distribution function $F(u) = P(U \leq u)$, we define $\xi(q) = \inf\{u : q \leq F(u)\}$ as the q -th quantile, and let $\mathbb{M}(U) = \xi(0.50)$ denote the median and $iqr(U) = \xi(0.75) - \xi(0.25)$ denote the interquartile range. Therefore, ψ_k^q captures the difference in median outcomes under treatment versus control, expressed in terms of the interquartile range under control. For example, $\psi_k^q = 0.5$ means treatment increases the median outcome, by half the interquartile range under control.

As for the scaled average effect in (1), a doubly robust and locally efficient estimator for the quantile effect ψ_k^q is given by

$$\hat{\psi}_k^q = \frac{\hat{F}_{1k}^{-1}(0.50) - \hat{F}_{0k}^{-1}(0.50)}{\hat{F}_{0k}^{-1}(0.75) - \hat{F}_{0k}^{-1}(0.25)} \quad (8)$$

where $\hat{F}_{ak}^{-1}(\cdot)$ is the inverse of $\hat{F}_{ak}(y) = \mathbb{P}_n\{\phi_{ak}^{(y)}(\mathbf{Z}; \hat{\pi}, \hat{v})\}$ (up to order $o_{\mathbb{P}}(1/\sqrt{n})$ if exact solutions cannot be found), for

$$\phi_{ak}^{(y)}(\mathbf{Z}; \pi, v) = \frac{\mathbb{1}(A = a)}{\pi(a|\mathbf{X})} \{ \mathbb{1}(Y_k \leq y) - v_k(y|\mathbf{X}, a) \} + v_k(y|\mathbf{X}, a) \quad (9)$$

where $v_k(y|\mathbf{X}, a) = \mathbb{P}(Y_k \leq y|\mathbf{X} = \mathbf{x}, A = a)$. We give the efficient influence function for ψ_k^q (as well as conditions under which this is the influence function for $\hat{\psi}_k^q$) in Section 4 of the Supplementary Materials; results for unscaled effects were developed by Díaz.²⁵

3.3 Hypothesis testing

In what follows, we suppose for concreteness that the mean-variance standardization in equation (1) is appropriate for all K outcomes. However, all results can be equally extended to the quantile-based standardization given in equation (7) (if mean-variance standardization is only appropriate for some covariates, we suggest using quantile-based standardization for all).

As mentioned in the Introduction, in studies with multiple outcomes it is often of interest to assess whether any outcomes are differentially affected by treatment, and if so, which outcomes. This can be particularly important when outcomes are meant to measure the same underlying construct. For example, Cohen et al.²⁶ and Roy et al.¹² studied effects of highly active antiretroviral therapy on several different measures of overall neurocognitive functioning among women with HIV. Additional motivation is given in Section 5, as well as by Lin et al.¹¹ and Roy et al.¹² This goal can be accomplished by testing a hypothesis of the form

$$H_0: \psi_1 = \psi_2 = \cdots = \psi_K \quad (10)$$

which says that all scaled treatment effects are equal. Recall that this is a meaningful hypothesis even if outcomes are measured on different scales, due to the fact that the scaled effects are unitless after dividing by standard deviations under control.

A doubly robust test of the hypothesis in equation (10) can be constructed based on the asymptotic distribution of

$$T_n = n(\mathbf{C}\hat{\psi})^T (\mathbf{C}\hat{\Sigma}\mathbf{C}^T)^{-1} (\mathbf{C}\hat{\psi}), \quad (11)$$

where $\mathbf{C}$ is a $(K-1) \times K$ banded matrix with elements $C_{ij} = \mathbb{1}(i=j) - \mathbb{1}(i=j-1)$, $\hat{\boldsymbol{\psi}} = (\hat{\psi}_1, \dots, \hat{\psi}_K)^T$, and $\hat{\boldsymbol{\Sigma}}$ is an estimator of the asymptotic variance $\boldsymbol{\Sigma}$ of $\hat{\boldsymbol{\psi}}$ (estimation of which is discussed in the next paragraph). Specifically, under conditions given in Sections 2 and 5 of the online Supplementary Materials, we have that

$$T_n \rightsquigarrow \chi_{K-1}^2 \quad (12)$$

under the null hypothesis H_0 of homogeneous effects given in equation (10). Therefore, an asymptotic p -value for testing H_0 is given by $P(\chi_{K-1}^2 \geq t_n)$, where t_n is the observed value of T_n in the sample (and χ_{K-1}^2 is a chi-squared random variable with $K-1$ degrees of freedom).

If the conditions given in Sections 2 and 5 of the online Supplementary Materials hold, then a closed-form estimator for $\boldsymbol{\Sigma}$ can be obtained by replacing unknown quantities in equation (6) with estimates and computing the empirical covariance

$$\hat{\boldsymbol{\Sigma}} = \mathbb{P}_n \{ \boldsymbol{\varphi}(\mathbf{Z}; \hat{\pi}, \hat{\mu}, \hat{\eta})^{\otimes 2} \}$$

where $\boldsymbol{\varphi} = (\varphi_1, \dots, \varphi_K)^T$ is a vector of the stacked influence functions from equation (6) for $k = 1, \dots, K$, and $\mathbf{u}^{\otimes 2} = \mathbf{u}\mathbf{u}^T$ for any vector $\mathbf{u}$. Note that the estimated influence functions used to construct the estimator will also depend on estimates of $\hat{\mathbb{E}}\{(Y_k^0)^2\} = \mathbb{P}_n\{\phi_{2k}(\mathbf{Z}; \hat{\pi}, \hat{\eta})\}$, $\hat{\mathbb{E}}(Y_k^0) = \mathbb{P}_n\{\phi_{0k}(\mathbf{Z}; \hat{\pi}, \hat{\mu})\}$, and $\hat{sd}(Y_k^0)^2 = \hat{\mathbb{E}}\{(Y_k^0)^2\} - \{\hat{\mathbb{E}}(Y_k^0)\}^2$.

The conditions for the above estimator to be valid require that the product of convergence rates for $\hat{\pi}$ and $(\hat{\mu}, \hat{\eta})$ is faster than $\sqrt{n}$, for example, if $\hat{\pi}$ is estimated with a correct parametric model or known and $(\hat{\mu}, \hat{\eta})$ is merely consistent (so that the product is $O_{\mathbb{P}}(1/\sqrt{n})o_{\mathbb{P}}(1) = o_{\mathbb{P}}(1/\sqrt{n})$), or if $(\hat{\pi}, \hat{\mu}, \hat{\eta})$ are all consistent and converge at faster than $n^{1/4}$ rates (so that the product is $o_{\mathbb{P}}(n^{-1/4})o_{\mathbb{P}}(n^{-1/4}) = o_{\mathbb{P}}(1/\sqrt{n})$). There is one setting where the above approach is valid even if this condition on the product of convergence rates does not hold. Specifically, if $\hat{\pi}$ is estimated with a correct parametric model, then even if $(\hat{\mu}, \hat{\eta})$ is misspecified and the estimator $\hat{\boldsymbol{\Sigma}}$ is thus inconsistent, the above approach gives conservative p -values and is still valid. This is a result of the fact that estimating the propensity score π when it is actually known cannot decrease (and will generally increase) efficiency¹⁶; thus $\hat{\boldsymbol{\Sigma}} \geq \boldsymbol{\Sigma}$ in the sense that $\hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}$ is a positive definite matrix.

Alternatively the bootstrap can also be used to construct the estimator $\hat{\boldsymbol{\Sigma}}$; such an approach would be valid as long as $\hat{\boldsymbol{\psi}}$ is asymptotically linear, which is a weaker condition than requiring the product of convergence rates to be $o_{\mathbb{P}}(1/\sqrt{n})$. (For example, asymptotic linearity would hold in the scenarios discussed above, as well as if $(\hat{\pi}, \hat{\mu}, \hat{\eta})$ were estimated with parametric models and it was only assumed that either $\hat{\pi}$ or $(\hat{\mu}, \hat{\eta})$ were correctly modeled.) In

practice, the bootstrap might be preferred for computing $\hat{\Sigma}$ since it depends on weaker assumptions, although it is more computationally expensive.

To test which outcomes are differentially affected by treatment, we can test the pairwise hypotheses $H_{jk} : \psi_j = \psi_k$. To control the family-wise error rate (which might be reasonable if K is not too large) a simple Bonferroni correction could be used. Alternatively, to control the false discovery rate (which might be preferable if K is large), the Benjamini–Hochberg procedure could be used instead.

3.4 Effect modification

Often it is of interest to go beyond marginal effects like $\mathbb{E}(Y_k^1 - Y_k^0)$ or the scaled version in equation (1), and further assess how treatment effects vary with covariates. This can be useful for exploring the mechanism by which treatment actually works, as well as for learning how to tailor treatment decisions to individual patient characteristics (since treatments may only work for some patients, or may be harmful for some and beneficial for others).

A natural extension of the standard effect parameter in equation (1) that allows for assessing such effect modification is given by

$$\gamma_k(\mathbf{v}) = \frac{\mathbb{E}(Y_k^1 - Y_k^0 | \mathbf{V} = \mathbf{v})}{sd(Y_k^0 | \mathbf{V} = \mathbf{v})} \quad (13)$$

where $\mathbf{V} \subseteq \mathbf{X}$ is a subset of the full covariate set that only includes the variables for which effect modification is of interest. Similar to prior subsections, this effect measures the mean difference in outcomes under treatment versus control for those with covariates $\mathbf{V} = \mathbf{v}$, in terms of the standard deviation under control for this same group.

When $\mathbf{V}$ only contains a modest number of discrete variables, the estimator in equation (4) can be easily modified to estimate $\gamma_k(\mathbf{v})$ with

$$\frac{\mathbb{P}_n\{\phi_{1k}^{(\mathbf{v})}(\mathbf{Z}; \hat{\pi}, \hat{\mu}) - \phi_{0k}^{(\mathbf{v})}(\mathbf{Z}; \hat{\pi}, \hat{\mu})\}}{\sqrt{\mathbb{P}_n\{\phi_{2k}^{(\mathbf{v})}(\mathbf{Z}; \hat{\pi}, \hat{\eta})\} - [\mathbb{P}_n\{\phi_{0k}^{(\mathbf{v})}(\mathbf{Z}; \hat{\pi}, \hat{\mu})\}]^2}} \quad (14)$$

where the functions $\phi_{ak}^{(\mathbf{v})}(\mathbf{Z}; \pi, \mu) = \phi_{ak}(\mathbf{Z}; \pi, \mu) \mathbb{1}(\mathbf{V} = \mathbf{v}) / \mathbb{P}_n\{\mathbb{1}(\mathbf{V} = \mathbf{v})\}$ and

$\phi_{2k}^{(\mathbf{v})}(\mathbf{Z}; \pi, \mu) = \phi_{2k}(\mathbf{Z}; \pi, \mu) \mathbb{1}(\mathbf{V} = \mathbf{v}) / \mathbb{P}_n\{\mathbb{1}(\mathbf{V} = \mathbf{v})\}$ are $\mathbf{v}$ -specific versions of the influence functions from previous sections, so that the averages in equation (14) are just over those units with $\mathbf{V} = \mathbf{v}$. The influence function for the estimator in equation (14) is given in the next subsection.

When $\mathbf{V}$ contains a continuous variable or many discrete variables, the above approach will not be feasible since the cells $\mathbf{V} = \mathbf{v}$ will be very small or empty. There are a few options for

such cases. First, one could estimate $\mathbb{E}(Y_k^a|\mathbf{V})$ and $\mathbb{E}\{(Y_k^a)^2|\mathbf{V}\}$ by regressing $\phi_{ak}(\mathbf{Z}; \hat{\pi}, \hat{\mu})$ and $\phi_{2k}(\mathbf{Z}; \hat{\pi}, \hat{\eta})$ on $\mathbf{V}$ using any preferred methods, such as parametric regression modeling or flexible machine learning, and then construct

$$\hat{\gamma}_k(\mathbf{v}) = \frac{\hat{\mathbb{E}}(Y_k^1|\mathbf{V}=\mathbf{v}) - \hat{\mathbb{E}}(Y_k^0|\mathbf{V}=\mathbf{v})}{\sqrt{\hat{\mathbb{E}}\{(Y_k^0)^2|\mathbf{V}=\mathbf{v}\} - \{\hat{\mathbb{E}}(Y_k^0|\mathbf{V}=\mathbf{v})\}^2}}$$

based on these regressions. A potential disadvantage of this approach is that $\hat{\gamma}_k(\mathbf{v})$ will in general not follow an interpretable parametric model, even if the models for the components $\mathbb{E}(Y_k^a|\mathbf{V})$ and do; for example, the ratio of two linear models is not itself linear.

To combat this problem, one could regress the predicted values from the above approach onto a parametric model $\gamma_k(\mathbf{v}; \boldsymbol{\theta})$ for some $\boldsymbol{\theta} \in \mathbb{R}^q$ (e.g., the linear model $\gamma_k(\mathbf{v}; \boldsymbol{\theta}) = \boldsymbol{\theta}^T \mathbf{v}$); here $\mathbf{v}$ can be parametrized so that the dimension q is its cardinality. Alternatively, semiparametric estimators could be developed under the restriction $\gamma_k(\mathbf{v}) = \gamma_k(\mathbf{v}; \boldsymbol{\theta})$; we leave this to future work. Lastly, another alternative, which is simple and easy to implement, would be to standardize by the marginal standard deviation $sd(Y_k^0)$ instead of the conditional form $sd(Y_k^0|\mathbf{V})$; this could be justified merely as an alternative standardization or via the assumption that $sd(Y_k^0|\mathbf{V}) = sd(Y_k^0)$.

3.5 Weighted average summary measures

So far we have discussed robust estimation of outcome-specific effects (scaled for comparability) in Sections 3.1, 3.2, and 3.4, as well as testing of homogeneous effect hypotheses in Section 3.3. In many studies, it may also be useful to report a summary measure of treatment effect across outcomes. For this purpose, we propose weighted average effects of the form

$$\psi^* = \sum_{k=1}^K \mathbb{E} \left\{ w_k(\mathbf{V}) \frac{\mathbb{E}(Y_k^1 - Y_k^0|\mathbf{V})}{sd(Y_k^0|\mathbf{V})} \right\} = \sum_{k=1}^K \mathbb{E} \{ w_k(\mathbf{V}) \gamma_k(\mathbf{V}) \}, \quad (15)$$

where $w_k(\mathbf{v})$ is an arbitrary user-specified function that indicates how much the summary should be weighted toward outcome k and strata $\mathbf{V} = \mathbf{v}$. For example, to weight outcomes equally and to weight strata according to the marginal distribution of $\mathbf{V}$, one could use $w_k(\mathbf{v}) = 1$; if some outcomes or strata were (a priori) more important, weights could be adjusted accordingly.

An interesting example where weighted average summary measures like ψ^* are useful is as follows. In Kangovi et al.,² a follow-up study to Kangovi et al.¹ patients were asked (prior to randomization) which of K outcomes they would like to focus on for improvement, so that $V \in 2\{1, \dots, K\}$ indicates this choice. Investigators were interested in the overall average

effect of treatment A on patients' selected outcomes, which is a special case of the weighted summary in equation (15). Specifically, by selecting $w_k(V) = \mathbb{1}(V = k)$ the above summary measure reduces to

$$\sum_{k=1}^K \mathbb{P}(V = k) \left\{ \frac{\mathbb{E}(Y_k^1 - Y_k^0 | V = k)}{sd(Y_k^0 | V = k)} \right\}.$$

This is a weighted average of the scaled effects $\gamma_k(k)$ on selected outcomes (among patients selecting these outcomes), where the weights equal the proportion of patients choosing to focus on each outcome. It should be noted, of course, that these summary measures can provide an obscured view of the multivariate effects $\gamma_k(\mathbf{v})$, $k = 1, \dots, K$, when they are heterogeneous; thus in practice such summary measures should be presented alongside estimates of outcome-specific effects.

The efficient influence function for the parameter ψ^* in equation (15) is given by

$$\sum_{k=1}^K w_k(\mathbf{V}) \{ \varphi_k^{(\mathbf{v})}(\mathbf{Z}) + \gamma_k(\mathbf{V}) \} - \psi^* \quad (16)$$

where $\varphi_k^{(\mathbf{v})}(\mathbf{Z})$ is the efficient influence function for the $\mathbf{v}$ -specific effect $\gamma_k(\mathbf{v})$ given by

$$\frac{\mathbb{1}(\mathbf{V} = \mathbf{v})}{\mathbb{P}(\mathbf{V} = \mathbf{v})} \left(\frac{\phi_{1k}(\mathbf{Z}) - \phi_{0k}(\mathbf{Z})}{sd(Y_k^0 | \mathbf{V})} - \gamma_k(\mathbf{v}) \left[\frac{\phi_{2k}(\mathbf{Z}) + \mathbb{E} \left\{ (Y_k^0)^2 | \mathbf{V} \right\} - 2\mathbb{E}(Y_k^0 | \mathbf{V})\phi_{0k}(\mathbf{Z})}{2sd(Y_k^0 | \mathbf{V})^2} \right] \right)$$

Therefore, as in previous sections, Wald-type confidence intervals for ψ^* can be constructed by estimating the values of the influence function in equation (16) and using the corresponding empirical standard error.

4 Simulation study

To illustrate some of our proposed methods and explore finite-sample performance, we simulated data with $K = 4$ outcomes from the following model

$$\mathbf{X} \sim N(0, I_4).$$

$$A | \mathbf{X} \sim \text{Bernoulli} \{ \pi(\mathbf{x}) \},$$

$$\pi(\mathbf{x}) = \text{expit} \{ (2x_1 - 4x_2 + 2x_3 - x_4)/4 \},$$

$$Y_k \mid \mathbf{X}, A \sim N\{\mu_k(\mathbf{x}, a), k^2\},$$

$$\mu_k(\mathbf{x}, a) = k \sum_{j \neq k} (-1)^{j+k-1} x_j + 2(k-\lambda)a$$

For the main setting, we consider (where $\lambda = 2$), the above model gives true values of $\mathbb{E}(Y_k^1 - Y_k^0) = 2(k-2)$ and $sd(Y_k^0) = 2k$, so that $\psi_k = 1 - 2/k$, that is

$$\boldsymbol{\psi} = (-1, 0, 1/3, 1/2)^T.$$

The above model also implies that that $\eta_k(\mathbf{x}, a)$ follows a linear model that is quadratic in the covariates and includes all two-way interactions.

To analyze the above simulated data, we considered estimation of the scaled effect parameter $\boldsymbol{\psi}$, as well as testing of the homogeneous effects hypothesis discussed in Section 3.3. Here, we implemented the proposed estimating equation version of our estimator $\hat{\boldsymbol{\psi}}$, but also give results for the TMLE version in Section 7 of the online Supplementary Materials (with R code for both given in Section 6). This estimator depends on estimates of the nuisance functions $(\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\eta})$, which we constructed using correctly specified parametric models. We used parametric models to ease computation and focus ideas, but in practice we suggest using more flexible methods to minimize risk of model misspecification. To assess potential impacts of such model misspecification in our simulation setup, we fit parametric models with transformed versions of the covariates, using the same transformations as Kang and Schafer.²⁷ Confidence intervals were constructed using the closed-form influence function-based approach, which is technically only valid under correct modeling of $\boldsymbol{\pi}$; in practice the bootstrap could be used if misspecification of $\boldsymbol{\pi}$ is possible. Results for estimating $\boldsymbol{\psi}$ are given in Table 1 (with RMSE scaled by $\sqrt{n}$ for easier interpretation).

The simulation results reflect what is expected based on theory. In particular, the scaled effect $\boldsymbol{\psi}$ was estimated with small bias whenever either $\boldsymbol{\pi}$ or $(\boldsymbol{\mu}, \boldsymbol{\eta})$, were correctly modeled, indicating the double robustness of our approach. Of course, if all nuisance estimators are misspecified, no method can promise small bias. Finite-sample biases under correct specification of either $\boldsymbol{\pi}$ or $(\boldsymbol{\mu}, \boldsymbol{\eta})$ were quite small even when $n = 200$, and for $n = 1000$ they were almost always zero after rounding (i.e., less than 0.0005). When all nuisance estimators were correctly modeled, coverage was very close to the nominal 95% level, especially for the $n = 1000$ case. Under misspecification of either $\boldsymbol{\pi}$ or $(\boldsymbol{\mu}, \boldsymbol{\eta})$, coverage was usually close to 95% even without any theoretical guarantees, while (as expected) coverage was poor under complete misspecification. In cases where one of $\boldsymbol{\pi}$ or $(\boldsymbol{\mu}, \boldsymbol{\eta})$ was incorrectly modeled, we expect the bootstrap to provide improved performance. Results for the TMLE estimator were similar (see Section 7 of the online Supplementary Materials).

For the hypothesis testing portion of the simulation study, we used the proposed test statistic T_n based on the estimators $\hat{\psi}$ described earlier, with the covariance estimator $\hat{\Sigma}$ proposed in Section 3.3 (the estimated covariance of the estimated influence function values). We implemented the approach under the null setting $\lambda = 0$ to assess type I error (when $\lambda = 2$ and the homogeneous effect hypothesis fails to hold, our approach gave 100% power in all settings). As before, our theory only guarantees correct error control under correct model specification for both μ and (π, η) when using the closed-form variance estimator.

The simulations show that our test approximately controls type I error at nominal rates in large samples, if either working model is correct. There is some anticonservative bias for smaller sample sizes, as is often the case for generalized Wald tests.²⁸ This might be ameliorated with the bootstrap, either to estimate the variance or if it could be possible to resample under the null directly; we hope to pursue this in future work. Although the closed-form variance estimator does not guarantee error control under misspecification, the type I error was relatively close to 5% as long as one of π or (μ, η) was correctly modeled. Again TMLE results are given in Section 7 of the online Supplementary Materials; in this setting the TMLE version of our estimator gave inflated type I error relative to the estimating equation results presented in the main text.

5 Application

Here, we apply our proposed methods to a recent trial² studying the effects of a (randomized) community health worker intervention on four chronic disease-related outcomes: cigarettes per day (CPD), systolic blood pressure (SBP), HbA1c, and body mass index (BMI). At baseline each of $n = 302$ subjects was asked which of the outcomes he or she would prefer to focus on, and the primary goal of the study was to learn whether the intervention affected patients' focus outcomes. Each outcome was measured in terms of change between study enrollment and 6-month follow-up, so negative values indicate improvement. Kangovi et al.^{1,2,29} give full details of the study population and design; we give summary statistics and histograms of the outcome distributions in the online Supplementary Materials.

In our analysis, we adjust for baseline outcomes, age, gender, and the selected focus outcome, and we used the cross validation-based Super Learner²⁰ to combine parametric models (linear for the outcome, logistic for the treatment), generalized additive models, and random forests. Note however that here covariate adjustment is only for the purposes of increasing efficiency, since the intervention was completely randomized (i.e., our propensity score model is guaranteed to be correctly specified). We estimate the scaled effects from Section 3.1, test effect homogeneity as in Section 3.3, estimate how effects vary with $V =$ selected outcomes as in Section 3.4, and finally estimate the average effect on selected outcomes via the summary measure approach from Section 3.5. We use mean-variance scaling since our outcomes have symmetric distributions, and because there is some discreteness in the outcome measurements so quantiles are not always uniquely defined.

Results are displayed in Figure 1. The scaled effect estimates (with nonsimultaneous 95% confidence intervals) were -0.22 ($-0.36, -0.08$) for CPD, 0.09 ($-0.04, 0.22$) for SBP, -0.08

(−0.22, 0.05) for HbA1c, and 0.02 (−0.15, 0.19) for BMI. These results indicate that the intervention was effective in reducing cigarettes smoked per day, yielding a decrease of roughly a fifth of a standard deviation (of pre-post differences). For reference, we give summary statistics for each outcome in the online Supplementary Materials. The null hypothesis of zero scaled effect for CPD was in fact rejected at level 0.05 after Bonferroni correction for multiple testing (since $p = 0.001 < 0.05/4$). There was some effect homogeneity, as indicated in the left panel of Figure 1 and by the fact that the homogeneous effects test from Section 3.3 was rejected ($p = 0.031$). Hence, we have evidence that the intervention affects different outcomes differently. Finally, we also see that effects were generally stronger for outcomes that patients selected to focus on. However, it is still only the CPD outcome for which we can reject a non-zero effect at the 0.05 level ($p = 0.003$). The estimated average effect on selected outcomes was a tenth of a standard deviation, and significantly different from zero. However, it is clear that much of this effect comes from the effects on smoking.

6 Discussion

In this paper, we proposed flexible methods for estimating and testing effects in studies with multiple outcomes. We developed nonparametric doubly robust methods for estimating effects scaled using mean-variance and interquartile range standardizations, including effect modification and weighted summary measures, and constructed a test of effect homogeneity. We expect our work to be important for both randomized trials and observational studies, and feel our distribution-free results fill an important gap in the literature.

There are a number of important future directions to this research. In an upcoming paper, we will apply the methods to explore effects of a community health worker intervention in a patient-centered study that allows patients to choose to focus on certain outcomes rather than others, as briefly described in Section 3.5. It will also be useful to develop tests of other hypotheses beyond homogeneity (e.g., tests of no effect), to more thoroughly explore the implementation and performance of the bootstrap or other methods for weakening assumptions for valid inference, and to consider an asymptotic regime in which $K = K_n$ increases with sample size (which may more accurately represent studies with many outcomes).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Emin Tahirovic for helpful discussions and programming support on an earlier version of this manuscript.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Edward Kennedy gratefully acknowledges support from NIH grant (R01-DK090385) and Shreya Kangovi and Nandita Mitra from PCORI grant (AD-1310-0792).

References

1. Kangovi S, Mitra N, Grande D, et al. Patient-centered community health worker inter-vention to improve post-hospital outcomes: a randomized clinical trial. *JAMA Int Med.* 2014; 174:535–543.
2. Kangovi S, Mitra N, Turr L, et al. A randomized controlled trial of a community health worker intervention in a population of patients with multiple chronic diseases: study design and protocol. *Contemp Clin Trials.* 2017; 53:115–121. [PubMed: 27965180]
3. VanderWeele TJ. Outcome-wide epidemiology. *Epidemiology.* 2017; 28:399–402. [PubMed: 28166102]
4. O'Brien PC. Procedures for comparing samples with multiple endpoints. *Biometrics.* 1984; 40:1079–1087. [PubMed: 6534410]
5. Pocock SJ, Geller NL, Tsiatis AA. The analysis of multiple endpoints in clinical trials. *Biometrics.* 1987; 43:487–498. [PubMed: 3663814]
6. Sammel M, Lin X, Ryan LM. Multivariate linear mixed models for multiple outcomes. *Stat Med.* 1999; 18:2479–2492. [PubMed: 10474154]
7. Freemantle N, Calvert M, Wood J, et al. Composite outcomes in randomized trials: greater precision but with greater uncertainty? *J Am Med Assoc.* 2003; 289:2554–2559.
8. Thurston SW, Ruppert D, Davidson PW. Bayesian models for multiple outcomes nested in domains. *Biometrics.* 2009; 65:1078–1086. [PubMed: 19397588]
9. Teixeira-Pinto A, Mauri L. Statistical analysis of noncommensurate multiple outcomes. *Circ Cardiovasc Qual.* 2011; 4:650–656.
10. Yoon FB, Fitzmaurice GM, Lipsitz SR, et al. Alternative methods for testing treatment effects on the basis of multiple outcomes: simulation and case study. *Stat Med.* 2011; 30:1917–1932. [PubMed: 21538986]
11. Lin X, Ryan LM, Sammel M. A scaled linear mixed model for multiple outcomes. *Biometrics.* 2000; 56:593–601. [PubMed: 10877322]
12. Roy JA, Lin X, Ryan LM. Scaled marginal models for multiple continuous outcomes. *Biostatistics.* 2003; 4:371–383. [PubMed: 12925505]
13. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *J Educ Psychol.* 1974; 66:688.
14. Bickel, PJ, Klaassen, CA, Ritov, Y, et al. Efficient and adaptive estimation for semiparametric models. Johns Hopkins University Press; 1993.
15. van der Laan, MJ, Robins, JM. Unified methods for censored longitudinal data and causality. New York, NY: Springer; 2003.
16. Tsiatis, AA. Semiparametric theory and missing data. New York, NY: Springer; 2006.
17. Kennedy, EH. Semiparametric theory and empirical processes in causal inference. In: He, Hua; Wu, Pan; Chen, Ding-Geng (DIN), editors. Statistical causal inferences and their applications in public health research. New York, NY: Springer; 2016. 141–167.
18. Robins JM, Rotnitzky A. Comments on “Inference for semiparametric models: some questions and an answer”. *Stat Sin.* 2001; 11:920–936.
19. Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. *Biometrics.* 2005; 61:962–973. [PubMed: 16401269]
20. van der Laan, MJ, Rose, S. Targeted learning: causal inference for observational and experimental data. New York, NY: Springer; 2011.
21. van der Vaart AW. Higher order tangent spaces and influence functions. *Stat Sci.* 2014; 29:679–686.
22. van der Vaart, AW, Wellner, JA. Weak convergence and empirical processes. New York, NY: Springer; 1996.
23. van der Vaart, AW. Asymptotic statistics. Cambridge, UK: Cambridge University Press; 2000.
24. van der Laan MJ, Rubin DB. Targeted maximum likelihood learning. UC Berkeley Division of Biostatistics Working Paper Series. 2006; 212:1–87.
25. Diaz I. Efficient estimation of quantiles in missing data models. arXiv.

26. Cohen RA, Boland R, Tashima KT, et al. Neurocognitive performance enhanced by highly active antiretroviral therapy in HIV-infected women. *AIDS*. 2001; 15:341–345. [PubMed: 11273214]
27. Kang JD, Schafer JL. Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Stat Sci*. 2007; 22:523–539.
28. Boos, DD, Stefanski, LA. *Essential statistical inference: theory and methods*. Springer; 2013.
29. Kangovi S, Mitra N, Grande D, et al. Community health worker support for disadvantaged patients with multiple chronic diseases: a randomized clinical trial. *Am J Public Health*. 2017; 107:1660–1667. [PubMed: 28817334]

Figure 1.

Estimates of scaled effects ψ_k (left) and effects on focus outcomes $\gamma_k(k)$ (right), with pointwise 95% confidence interval. Also shown is the p -value for testing effect homogeneity (left) and the estimated summary effect on selected outcomes with 95% confidence interval (right).

Table 1Results for estimating ψ across 1000 simulations.

Correct model and parameter	$n = 200$					$n = 1000$				
	Bias	SE	RMSE	Cov (%)		Bias	SE	RMSE	Cov (%)	
Both										
ψ_1	-0.02	0.14	2.03	90.2		-0.00	0.06	1.97	93.5	
ψ_2	0.00	0.09	1.32	93.0		0.00	0.04	1.25	95.6	
ψ_3	-0.00	0.09	1.27	94.9		-0.00	0.04	1.21	94.3	
ψ_4	0.02	0.11	1.54	92.8		0.00	0.05	1.52	93.9	
Trt										
ψ_1	-0.05	0.16	2.67	93.7		-0.01	0.08	2.63	96.8	
ψ_2	0.02	0.14	1.96	97.5		0.00	0.06	1.88	98.3	
ψ_3	-0.01	0.09	1.29	93.8		-0.00	0.04	1.32	95.6	
ψ_4	0.03	0.16	2.54	95.9		0.01	0.07	2.35	97.2	
Out										
ψ_1	-0.02	0.14	1.97	88.1		-0.00	0.06	1.74	93.0	
ψ_2	0.01	0.09	1.34	93.6		0.00	0.04	1.23	95.3	
ψ_3	0.00	0.09	1.29	93.3		-0.00	0.04	1.16	95.0	
ψ_4	0.01	0.11	1.60	89.2		0.01	0.05	1.48	91.8	
None										
ψ_1	-0.31	0.20	5.24	56.5		-0.30	0.09	9.90	5.8	
ψ_2	0.22	0.15	3.79	59.2		0.23	0.07	7.56	6.6	
ψ_3	-0.04	0.09	1.43	92.2		-0.03	0.07	2.59	87.6	
ψ_4	0.36	0.21	5.93	35.9		0.37	0.09	11.98	1.7	

Table 2

Results for testing homogeneity across 1000 simulations.

Correct model	Type I error (%)			
	$n = 200$	$n = 500$	$n = 1000$	$n = 5000$
Both	11.4	8.4	5.8	5.5
Trt	9.0	6.4	5.6	3.8
Out	13.5	9.3	7.7	7.5
None	55.7	91.2	99.2	100