



HHS Public Access

Author manuscript

Stat Med. Author manuscript; available in PMC 2022 August 30.

Published in final edited form as:

Stat Med. 2021 August 30; 40(19): 4310–4326. doi:10.1002/sim.9031.

Transporting Experimental Results with Entropy Balancing

Kevin P. Josey¹, Seth A. Berkowitz², Debashis Ghosh¹, Sridharan Raghavan^{*,3,4}

¹Department of Biostatistics and Informatics, Colorado School of Public Health, University of Colorado Anschutz Medical Campus, Colorado, USA

²Division of General Medicine and Clinical Epidemiology, University of North Carolina at Chapel Hill School of Medicine, North Carolina, USA

³Rocky Mountain Regional VA Medical Center, Colorado, USA

⁴Division of Hospital Medicine, University of Colorado School of Medicine, Colorado, USA

Summary

We show how entropy balancing can be used for transporting experimental treatment effects from a trial population onto a target population. This method is doubly-robust in the sense that if either the outcome model or the probability of trial participation is correctly specified, then the estimate of the target population average treatment effect is consistent. Furthermore, we only require the sample moments of the effect modifiers drawn from the target population to consistently estimate the target population average treatment effect. We compared the finite-sample performance of entropy balancing with several alternative methods for transporting treatment effects between populations. Entropy balancing techniques are efficient and robust to violations of model misspecification. We also examine the results of our proposed method in an applied analysis of the Action to Control Cardiovascular Risk in Diabetes Blood Pressure (ACCORD-BP) trial transported to a sample of US adults with diabetes taken from the National Health and Nutrition Examination Survey (NHANES) cohort.

Keywords

Calibration; Causal Inference; Generalizability; Effect Modification

1 | INTRODUCTION

In a randomized controlled trial (RCT), the population from which the sample is collected, the *trial population*, often differs from the population of interest, the *target population*. This scenario becomes problematic when the true causal effect is heterogeneous, implying the existence of effect modifying covariates - *effect modifiers* - which alter the average

*Correspondence: Sridharan Raghavan, Rocky Mountain Regional VA Medical Center, sridharan.raghavan@va.gov.

Present Address: 1700 N Wheeling St, Aurora, CO 80045

Publisher's Disclaimer: Disclaimer: This manuscript will be submitted to the Department of Biostatistics and Informatics in the Colorado School of Public Health, University of Colorado Anschutz Medical Campus, in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Biostatistics for Kevin P. Josey. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or the United States Department of Veterans Affairs.

treatment effect. If the distribution of the effect modifiers is different in the trial and target populations, the average treatment effect observed in the trial will likely differ from what would be observed within the target population, limiting the conclusions that can be drawn from an otherwise well-designed study. It is worth noting that effect modifiers are specific to the scale of the target estimand. Throughout, we will refer to effect modifiers as variables that modify the average treatment effect which we will define later on.

The recent literature on the subject of transportability is divided into two scenarios determined by the nature of the trial and target populations, and the desired causal estimand. If the trial population is nested within the target population, we can extend the results of an RCT using a sample from the target population in a process called *generalizability*. If the target and trial populations are subpopulations drawn from some super population, then the problem is one of *transportability*. Figure 1 provides a diagram relating the data to the corresponding populations in both the generalizability and transportability problems. Note that for generalizability, the trial population is a subpopulation of the target population, while in transportability the target and trial populations are not nested. We will discuss the difference between these two scenarios in more detail in Section 2.2. The work herein, however, will focus primarily on the issue of transportability.

Some articles have approached the problem of transportability from the setting in which the investigator is provided the individual-level data from the trial population along with individual-level covariate data from the target population.¹ Another setting provides the individual-level data from the trial population, but only the covariate sample moments (e.g., the mean and standard deviation) from the target population, which can often be found in a so-called Table 1 throughout the medical literature.² One property that is often sought while developing estimators for causal inference is called double-robustness.³ In the context of transporting experimental results, this means that if either the probability of trial participation or the outcome model are correctly specified, then the resulting average treatment effect estimator is consistent.

We propose using entropy balancing to solve transportability problems. The procedure we propose builds upon several other causal effect estimators which employ convex optimization techniques to estimate a vector of *sampling weights*.^{2,4,5,6} These sampling weights would otherwise be uniform if the RCT data were randomly sampled from the target population. The literature on convex optimization in the context of causal inference has abounded in recent years.^{7,8,9} Rather than using these methods to exactly balance the covariate distribution between the treated and control units within an observational study, convex optimization techniques applied to transportability can be used to estimate weights which balance the covariate distribution of the trial participants and non-participants. Entropy balancing is flexible in that it can be applied both when the complete individual-level covariate data are provided as well as when only the covariate sample moments of the target population are provided, such as what might appear in the Table 1 of a clinical paper. Furthermore, the specific entropy balancing procedure we develop can be shown to be doubly-robust for estimating the target population average treatment effect given the complete individual-level covariate data in the context of transportability.

The entropy balancing solution we propose is also considered as a solution for indirectly comparing experimental results from two separate randomized trials with an anchored treatment arm, a problem not too dissimilar from that of transportability. However, transportability as we describe it does not require a second randomized trial. The sample drawn from the target population, which is subsequently used in our transportability formulations, does not require any information about the outcome or treatment assignment. Moreover, based on new results identified in the indirect comparison literature, we can relax a rather strong assumption about the nature of the potential outcomes typically made in the transportability literature.⁶ In addition, we take a more comprehensive view of entropy balancing through the lens of causal inference, motivating this work through the potential outcomes framework and describe several properties about entropy balancing for transportability that are sought for estimators in the causal inference literature. This includes the property of double-robustness, semiparametric efficiency, and considerations between finite-sample and superpopulation settings. We also compare the entropy balancing approach with several other methods in an effort to showcase the strengths of doubly-robust estimators more generally in transportability problems.

The contents of the article are as follows. In Section 2 we define the notation, setting, and assumptions necessary for transporting experimental results between populations and describe several existing methods for transportation, including two methods that can be applied in the setting where we are given only the sample covariate moments of the target population and two methods that require individual-level covariate data from the target population, one of which is doubly-robust. In Section 3, we introduce entropy balancing and describe the difference between conducting inference upon the target population average treatment effect versus the target sample average treatment effect. In Section 4 we compare the five methods considered in Sections 2 and 3 using numerical studies. We also illustrate through a secondary simulation how entropy balancing and other methods that do not require individual-level data from the target population only allow for inference upon the target sample average treatment effect and not the target population average treatment effect. In Section 5 we compare entropy balancing and inverse odds of sampling weights in a real-data example: transporting results from a clinical trial of blood pressure treatment intensity in diabetes patients to a representative sample of the US population. Section 6 concludes with a discussion.

2 | SETTING AND PRELIMINARIES

2.1 | Notation and Potential Outcomes

Suppose we have two random samples from different populations. For independent sampling units $i = 1, 2, \dots, n$, let $S_i \in \{0, 1\}$ denote a random sampling indicator. Indexed by $\{i : S_i = 1\}$, the *trial sample* evaluates the efficacy of some treatment on the trial population. The second sample is randomly selected from the target population and indexed by $\{i : S_i = 0\}$. We refer to this sample as the *target sample*. We denote $n_1 = \sum_{i=1}^n S_i$, $n_0 = \sum_{i=1}^n (1 - S_i)$, and $n = n_1 + n_0$. Both $\mathbb{E}(\cdot)$ and $\Pr\{\cdot\}$ will be evaluated over the *superpopulation* which is the combined trial and target population.

For $i = 1, 2, \dots, n$, let $\mathbf{X}_i \in \mathcal{X}$ denote a vector of measured baseline (i.e. pretreatment) covariates. For $i \in \{i: S_i = 1\}$, let $Y_i \in \mathfrak{R}$ denote the real valued outcome, and $Z_i \in \{0, 1\}$ denote the random treatment assignment. We assume throughout that \mathbf{X}_i contains an intercept term. The probability density function for \mathbf{X}_i is denoted $f(\mathbf{x}_i)$ for $\mathbf{x}_i \in \mathcal{X}$. Indexed by $j = 1, 2, \dots, m$, we denote the vector-valued balance function $\mathbf{c}(\mathbf{X}_i) \equiv [c_1(\mathbf{X}_i), c_2(\mathbf{X}_i), \dots, c_j(\mathbf{X}_i), \dots, c_m(\mathbf{X}_i)]^T$, which are the effect modifiers included into the models of S_i and Y_i along with Z_i . We suppose $c_1(\mathbf{X}_i) = 1$ for all $i = 1, 2, \dots, n$. Some examples for $c_j(\mathbf{X}_i)$, $j = 2, 3, \dots, m$ include polynomial transformations of the covariates and interaction terms - anything that might modify the effect of the treatment on the outcome. It is common practice to balance the covariates as well as the variance (i.e. second moments) of the covariates when more intimate knowledge about the effect modifying process is unknown.²

We employ the potential outcomes framework for a binary treatment.¹⁰ This framework allows us to construct the observed outcome in terms of the factual and counterfactual variables $Y_i(0)$ and $Y_i(1)$, $i = 1, 2, \dots, n$. $Y_i(0)$ and $Y_i(1)$ correspond to each unit's outcomes when $Z_i = 1$ and $Z_i = 0$, respectively. The observed responses are then defined as $Y_i \equiv Z_i Y_i(1) + (1 - Z_i) Y_i(0)$. The potential outcomes framework also allows us to define the target population average treatment effect, $\tau_{\text{TATE}} = \mathbb{E}[Y_i(1) - Y_i(0) | S_i = 0]$ and the target sample average treatment effect

$$\tau_{\text{TSATE}} \equiv \frac{1}{n_0} \sum_{\{i: S_i = 0\}} [Y_i(1) - Y_i(0)].$$

The target sample average treatment effect only concerns the effects among units within the target sample whereas the target population average treatment effect concerns the average effect for all units that make up the target population.¹¹ We also define $\rho(\mathbf{X}_i) \equiv \Pr\{S_i = 1 | \mathbf{X}_i\}$ and $\pi \equiv \Pr\{Z_i = 1 | \mathbf{X}_i\}$. Recall that in an RCT, $\pi \in (0, 1)$ should be independent of \mathbf{X}_i .

Another way to write τ_{TATE} is

$$\tau_{\text{TATE}} \equiv \frac{\mathbb{E}\{[1 - \rho(\mathbf{X}_i)][Y_i(1) - Y_i(0)]\}}{\mathbb{E}[1 - \rho(\mathbf{X}_i)]}.$$

This alternative definition identifies the target population average treatment effect as a type of weighted average treatment effect. A corollary to Theorem 4 of Hirano *et al.* can be used to derive the semiparametric efficiency bound for any estimator of τ_{TATE} as

$$\Sigma_{\text{semi}} \equiv \frac{\mathbb{E}\left\{[1 - \rho(\mathbf{X}_i)]^2 \left[\frac{\mathbb{V}[Y_i(1) | \mathbf{X}_i]}{\pi} + \frac{\mathbb{V}[Y_i(0) | \mathbf{X}_i]}{1 - \pi} + [\tau_{\text{TATE}}(\mathbf{X}_i) - \tau_{\text{TATE}}]^2 \right]\right\}}{\mathbb{E}[1 - \rho(\mathbf{X}_i)]^2} \tag{1}$$

where $\tau_{\text{TATE}}(\mathbf{X}_i) \equiv \mathbb{E}[Y_i(1) - Y_i(0) | \mathbf{X}_i, S_i = 0]$.¹² This setup allows us to utilize the asymptotic results about weighted average treatment effects derived previously using convex optimization techniques such as those employed by entropy balancing.¹³

We denote the population moments of the target covariate distribution as $\mathbb{E}[\mathbf{c}(\mathbf{X}_i) | S_i = 0] = \theta_0$. For much of this paper, we will describe methods for transporting experimental results which weight the responses \mathbf{Y}_i for $i \in \{i: S_i = 1\}$ so that the weighted trial sample moments are the same as the population moments of the target population.¹⁴ We will denote the sample weights as $\gamma \equiv (\gamma_1, \gamma_2, \dots, \gamma_{n_1})$. Since θ_0 is usually unknown, we will need to make use of the estimator $\hat{\theta}_0 \equiv n_0^{-1} \sum_{\{i: S_i = 0\}} \mathbf{c}(\mathbf{X}_i)$. In practice, we usually set $\mathbf{c}(\mathbf{X}_i) = \mathbf{X}_i$ unless more is known about the data generating mechanisms. In those cases, $\hat{\theta}_0$ typically appears in the so-called Table 1 of many publications.

2.2 | Assumptions for Transportability

The following assumptions facilitate our ability to transport experimental results onto a target population. These assumptions represent the necessary conditions required for the transporting experimental results. They are also adapted from similar articles on the subject.^{15,1,16} Furthermore, we invoke the stable unit treatment value, which comprises the no interference and consistency assumptions.

Assumption 1 (Mean Difference Exchangeability).—The target population average treatment effect conditional on the baseline covariates is exchangeable between samples:

$$\mathbb{E}[Y_i(1) - Y_i(0) | \mathbf{X}_i, S_i = 1] = \mathbb{E}[Y_i(1) - Y_i(0) | \mathbf{X}_i, S_i = 0] = \mathbb{E}[Y_i(1) - Y_i(0) | \mathbf{X}_i].$$

Assumption 2 (Sampling Positivity).—The probability of trial participation, conditioned on the baseline covariates necessary to ensure Assumption 1, is bounded away from zero and one:

$$0 < \Pr\{S_i = 1 | \mathbf{X}_i = x_i\} < 1 \text{ for all } x_i \in \mathcal{X} \text{ where } f(x_i | S_i = 0) > 0.$$

Assumption 3 (Strongly Ignorable Treatment Assignment).—The potential outcomes among the trial participants are independent of the treatment assignment given \mathbf{X}_i :

$$[Y_i(0), Y_i(1)]^T \perp\!\!\!\perp Z_i | \mathbf{X}_i \text{ for all } i \in \{i: S_i = 1\}.$$

Assumption 3 is a standard assumption in the potential outcomes literature.¹⁷ This assumption can be further simplified in an RCT setting to assume

$$[Y_i(0), Y_i(1)]^T \perp\!\!\!\perp Z_i \text{ for all } i \in \{i: S_i = 1\}$$

since there should be no association between the treatment assignment and the covariates. The covariate imbalance that requires amelioration in transportability instead appears between \mathbf{X}_i and S_i .

As noted previously in the Introduction, there are subtle distinctions between generalizability and transportability. The main difference occurs with the causal estimand of interest. In transportability, the target estimand is τ_{TATE} . For generalizability, the causal estimand of interest is $\tau_{\text{ATE}} = \mathbb{E}[Y_i(1) - Y_i(0)]$. This is on account of the trial population being nested within the target population, so the superpopulation and the target population are identical. Under our notation, generalizability further assumes that the units $\{i : S_i = 0\}$ are sampled from the target population and the complement of the trial population. As a result, we would need to rewrite Assumption 2 for generalizability to state

$$0 < \Pr\{S_i = 1 \mid \mathbf{X}_i = \mathbf{x}_i\} < 1 \text{ for all } \mathbf{x}_i \in \mathcal{X} \text{ where } f(\mathbf{x}_i) > 0.$$

We avoid this setup to the problem and instead focus on methods developed for transportability and inference on τ_{TATE} .

In addition to Assumptions 1–3, we require the following assumptions to establish the double-robustness property of entropy balancing. We will show that if either assumption is met, then the entropy balancing methods introduced in Section 3 will be consistent for the target population average treatment effect. We also use these assumptions to establish consistency for some of the other methods we describe in Section 2.3 when standard regression methods are employed. Note that an underlying requirement implied by these two assumptions is that there is no unmeasured effect modification present given the known balance functions. If an effect modifier is missing, then any estimator we present will likely produce biased estimates of the target population average treatment effect.

Assumption 4 (Conditional Linearity).—The expected value of the potential outcomes, conditioned on \mathbf{X}_i , is linear across the span of the covariates. That is $\mathbb{E}[Y_i(1) - Y_i(0) \mid \mathbf{X}_i] = \mathbf{c}(\mathbf{X}_i)^T \boldsymbol{\alpha}$ and $\mathbb{E}[Y_i(0) \mid \mathbf{X}_i] = \mathbf{c}(\mathbf{X}_i)^T \boldsymbol{\beta}$ for all $i = 1, 2, \dots, n$ and $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathfrak{R}^m$.

Assumption 5 (Linear Conditional Log-Odds).—The log-odds of trial participation are linear across the span of the covariates. That is $\text{logit}[\rho(\mathbf{X}_i)] = \mathbf{c}(\mathbf{X}_i)^T \boldsymbol{\lambda}$ for all $i = 1, 2, \dots, n$ and some $\boldsymbol{\lambda} \in \mathfrak{R}^m$.

Assumption 1 is substantially relaxed from what appears in more recent literature.^{1,16} These other articles require the expected value of the potential outcomes to be exchangeable between populations:

$$\begin{aligned} \mathbb{E}[Y_i(1) \mid \mathbf{X}_i, S_i = 1] &= \mathbb{E}[Y_i(1) \mid \mathbf{X}_i, S_i = 0] \quad \text{and} \quad \mathbb{E}[Y_i(0) \mid \mathbf{X}_i, S_i = 1] \\ &= \mathbb{E}[Y_i(0) \mid \mathbf{X}_i, S_i = 0]. \end{aligned} \tag{2}$$

The indirect comparison literature refers to this assumption as the conditional constancy of *absolute* effects whereas Assumption 1 is commonly referred to as the conditional constancy of *relative* effects. Whereas the conditional constancy of absolute effects requires adjustment for all prognostic and effect modifying covariates, the conditional constancy of relative effects only requires adjustments for the effect modifiers. Note that the much stronger conditional constancy of absolute effects assumption is enforced in Assumption 4. However, the less stringent Assumption 1, in addition to Assumptions 2 and 3, is all that is required to obtain consistent estimates when Assumption 5 holds. This assumption relaxation result can be made using arguments of anchored indirect comparisons.⁶ Suppose the target sample is a randomized control trial comparing $Z = 1$ with $Z = 0$, similar to the data contained within the trial sample, only we do not observe either the outcome or the treatment assignment. Then this target “trial” is figuratively anchored by both treatment groups. If the target “trial” were to be conducted, and the outcome and treatment data were collected, then a resulting indirect comparison estimator should yield estimates of zero as both “trials” examine the same estimand.¹⁸ This is precisely what transportability methods are targeting – what the causal effect would be if the trial were conducted on a different population. Only requiring the conditional constancy of relative effects assumption versus the conditional constancy of absolute effects assumption adds incentive to focus on correctly specifying the sampling model over the outcome model through the entropy balancing techniques that we will describe in Section 3.

2.3 | Alternative Methods for Transportability

In this section we present four different methods for transporting experimental results to estimate τ_{TATE} . For each method, we assume we are given Assumptions 1–3. The first method weights responses of the trial sample with the inverse odds of sampling.¹⁹ Define the inverse odds of sampling weights as

$$\hat{\gamma}_i^{\text{PS}} = \begin{cases} \frac{1 - \hat{\rho}(\mathbf{X}_i)}{\hat{\rho}(\mathbf{X}_i)\hat{\pi}}, & \text{when } S_i = 1, Z_i = 1 \\ \frac{1 - \hat{\rho}(\mathbf{X}_i)}{\hat{\rho}(\mathbf{X}_i)(1 - \hat{\pi})}, & \text{when } S_i = 1, Z_i = 0 \\ 0, & \text{when } S_i = 0 \end{cases}$$

where $\hat{\pi}$ is a consistent estimator of the probability of treatment and $\hat{\rho}(\mathbf{X}_i)$ is a consistent estimator of the probability of trial participation. The target population average treatment effect is then estimated by computing

$$\hat{\tau}_{\text{IOSW}} = \sum_{\{i: S_i = 1\}} \frac{\hat{\gamma}_i^{\text{PS}} Z_i Y_i}{\sum_{\{i: S_i = 1\}} \hat{\gamma}_i^{\text{PS}} Z_i} - \sum_{\{i: S_i = 1\}} \frac{\hat{\gamma}_i^{\text{PS}} (1 - Z_i) Y_i}{\sum_{\{i: S_i = 1\}} \hat{\gamma}_i^{\text{PS}} (1 - Z_i)}.$$

If Assumption 5 is given, we may use logistic regression to consistently estimate $\hat{\rho}(\mathbf{X}_i)$. A consistent estimator for $\rho(\mathbf{X}_i)$ by extension renders $\hat{\tau}_{\text{IOSW}}$ consistent for τ_{TATE} .

Given the extended conditional constancy of absolute effects assumption in (2), another proposed solution is to find a consistent estimator of the conditional means for the potential outcomes with the sample data; $\mu_1(\mathbf{X}_i) \equiv \mathbb{E}[Y_i(1)|\mathbf{X}_i, S_i = 1]$ and $\mu_0(\mathbf{X}_i) \equiv \mathbb{E}[Y_i(0)|\mathbf{X}_i, S_i = 1]$.

We will refer to this method as the outcome modeling (OM) approach. The consistent estimators are denoted as $\hat{\mu}_1(\mathbf{X}_i)$ and $\hat{\mu}_0(\mathbf{X}_i)$, respectively. Under Assumption 1, τ_{TATE} can be estimated by solving for

$$\hat{\tau}_{\text{OM}} = \frac{1}{n_0} \sum_{\{i: S_i = 0\}} [\hat{\mu}_1(\mathbf{X}_i) - \hat{\mu}_0(\mathbf{X}_i)].$$

In the causal inference literature, this method follows the framework for computing causal effects known as g-computation.²⁰ If Assumption 4 is given, we can estimate τ_{TATE} with the OM approach if we are only given $\hat{\theta}_0$ instead of \mathbf{X}_i for all $i \in \{i: S_i = 0\}$. To do so, we would regress Y_i onto $\mathbf{c}(\mathbf{X}_i)$ for all $\{i: S_i = 1, Z_i = 1\}$ and $\{i: S_i = 1, Z_i = 0\}$ to get $\hat{\alpha}$ and $\hat{\beta}$, respectively. We then compute

$$\hat{\tau}_{\text{OM}} = \hat{\theta}_0^T(\hat{\alpha} - \hat{\beta})$$

where $\hat{\alpha}$ and $\hat{\beta}$ can be fit with ordinary least squares.

The OM approach and the inverse odds of sampling weights may be combined into a so called doubly-robust estimator. A doubly robust (DR) estimator combines estimators of the model components, in this case the model for $[Y_i(1), Y_i(0)]$ and S_i as to be consistent when at most one model is misspecified. The conventional doubly-robust estimator for a binary treatment was first proposed as a semiparametric technique for missing-data problems.²¹ There have been extensive modifications to this conventional doubly-robust estimator, including alterations for transporting experimental results of a binary treatment.⁵ Using the same notation outlined for the outcome model approach and the inverse odds of sampling weights, the target population average treatment effect can be estimated by solving for

$$\begin{aligned} \hat{\tau}_{\text{DR}} = & \sum_{\{i: S_i = 1\}} \frac{\hat{\gamma}_i^{\text{PS}} Z_i [Y_i - \hat{\mu}_1(\mathbf{X}_i)]}{\sum_{\{i: S_i = 1\}} \hat{\gamma}_i^{\text{PS}} Z_i} - \sum_{\{i: S_i = 1\}} \frac{\hat{\gamma}_i^{\text{PS}} (1 - Z_i) [Y_i - \hat{\mu}_0(\mathbf{X}_i)]}{\sum_{\{i: S_i = 1\}} \hat{\gamma}_i^{\text{PS}} (1 - Z_i)} \\ & + \frac{1}{n_0} \sum_{\{i: S_i = 0\}} [\hat{\mu}_1(\mathbf{X}_i) - \hat{\mu}_0(\mathbf{X}_i)]. \end{aligned} \tag{3}$$

It is easy to see that if $\hat{\mu}_0(\mathbf{X}_i)$ and $\hat{\mu}_1(\mathbf{X}_i)$ are consistent for $\mu_0(\mathbf{X}_i)$ and $\mu_1(\mathbf{X}_i)$, respectively, then the last summation in (3) is consistent for the target sample average treatment effect (and therefore also for the target population average treatment effect) while the first two summations converge in probability to zero. Similarly, if $\hat{\gamma}^{\text{PS}}$ is consistent for

$\rho(\mathbf{X}_i)[1 - \rho(\mathbf{X}_i)]^{-1}$, then the first two summations of (3) will consistently estimate the negative bias induced by the last summation.¹⁶

Another doubly-robust estimator closely related to the augmented estimator in (3) uses targeted maximum likelihood estimation (TMLE).²² TMLE is a popular choice among causal practitioners due to its flexibility for estimating a variety of different estimands, including the target population average treatment effect.¹ For transportability, the targeted maximum likelihood estimator is constructed as follows. First, the initial estimates of $\hat{\mu}_1(\mathbf{X}_i)$ and $\hat{\mu}_0(\mathbf{X}_i)$ are fit using the trial sample data. We then update the predictions of the potential outcomes on the trial sample with

$$\begin{aligned} \tilde{\mu}_0(\mathbf{X}_i) &= \hat{\mu}_0(\mathbf{X}_i) + \hat{\epsilon}_0(1 - Z_i)\hat{\gamma}_i^{\text{PS}} \\ \tilde{\mu}_1(\mathbf{X}_i) &= \hat{\mu}_1(\mathbf{X}_i) + \hat{\epsilon}_1 Z_i \hat{\gamma}_i^{\text{PS}}. \end{aligned} \tag{4}$$

The estimates of ϵ_0 and ϵ_1 are obtained by regressing the outcome Y_i onto the clever covariates $-Z_i \hat{\gamma}_i^{\text{PS}}$ and $(1 - Z_i)\hat{\gamma}_i^{\text{PS}}$ – with $\hat{\mu}_0(\mathbf{X}_i)$ and $\hat{\mu}_1(\mathbf{X}_i)$ serving as offsets for all $i \in \{i : S_i = 1\}$. The estimator of τ_{TATE} under the TMLE framework solves for

$$\hat{\tau}_{\text{TMLE}} = \frac{1}{n_0} \sum_{\{i: S_i = 0\}} [\tilde{\mu}_1(\mathbf{X}_i) - \tilde{\mu}_0(\mathbf{X}_i)] \tag{5}$$

in a similar manner to the OM approach. Equation (5) is doubly-robust for estimating τ_{TATE} in the sense that if either the sampling model or the potential outcomes models are consistent, then $\hat{\tau}_{\text{TMLE}}$ is also consistent.¹ TMLE also requires individual-level covariate data to estimate some of the components in (4) and (5).

For the DR and TMLE methods, it is unclear to us whether the more relaxed Assumption 1 remains applicable in cases where the sampling model is correctly specified. Note that both these methods are heavily geared toward the outcome regression model being correctly specified. To avoid any distraction from this potential discrepancy, we ensure the conditional constancy of absolute effects assumption is satisfied in the simulation study found in Section 4.1. Furthermore, we will compare the entropy balancing methods described in the next section with IOSW alone in the data analysis found in Section 5 as the conditional constancy of absolute effects assumption cannot be guaranteed like in a simulation study.

3 | ENTROPY BALANCING

Entropy balancing has emerged as a popular method for estimating weights in a variety of contexts, particularly for estimating the average treatment effect of the treated.^{7,23} Entropy balancing has also previously been introduced for evaluating indirect comparisons of randomized trials, though in this case it is referred to as a method of moment estimator for the inverse odds of sampling when the probability of trial participation follows a logit model.² This method of moment estimator just so happens to be the dual solution to an

entropy balancing primal problem. We prefer to use the term entropy balancing as it more concisely describes the underlying constrained convex optimization problem that must be solved in order to balance the covariate distribution,

$$\begin{aligned} & \text{minimize} && \sum_{\{i: S_i = 1\}} \gamma_i \log \gamma_i - \gamma_i \\ & \text{subject to} && \sum_{\{i: S_i = 1\}} \mathbf{A}_i \gamma_i = \mathbf{B} \end{aligned} \tag{6}$$

Note that the criterion distance in (6) is the entropy function, hence the “entropy” in entropy balancing. The constraints in (6), represented by the vectors \mathbf{A}_i and \mathbf{B} , can be constructed to satisfy some moment balancing conditions, hence the “balancing” aspect of entropy balancing. For example we can set $\mathbf{A}_i = \mathbf{c}(\mathbf{X}_i)$ and $\mathbf{B} = \hat{\theta}_0$ so that the weighted sample moments of $\mathbf{c}(\mathbf{X}_i)$ for all $i \in \{i: S_i = 1\}$ are equal to the sample moments $\mathbf{c}(\mathbf{X}_i)$ for all $i \in \{i: S_i = 0\}$. This specific choice of \mathbf{A}_i and \mathbf{B} is the primal problem for the previously proposed method of moments estimator.² Entropy balancing and the method of moment estimator for evaluating indirect comparisons are often conflated due to the different Lagrangian dual solutions one can arrive at while solving (6), one of which we will get to later in this section. Nevertheless, due to the strict convexity of the criterion function, the solution to (6) is unique and hence the dual solution must also be unique.²⁴ This result was also made explicit specifically when the convex criterion function is the entropy function.¹⁸

The dual solution for the method of moments estimator only requires the target sample moments of the covariates. For this balancing solution, denote $\tilde{\mathbf{c}}(\mathbf{X}_i) = (2Z_i - 1, \mathbf{X}_i)$ and $\tilde{\theta}_0 = (0, \hat{\theta}_0^T)^T$. The method of moments estimator first solves the Lagrangian dual problem

$$\hat{\lambda} = \arg \max_{\lambda \in \mathfrak{R}^{m+1}} \sum_{\{i: S_i = 1\}} \left[-\exp(-\tilde{\mathbf{c}}(\mathbf{X}_i)^T \lambda) - \tilde{\theta}_0^T \lambda \right], \tag{7}$$

which in turn is used to estimate the sampling weights, $\hat{\gamma}_i^{\text{MOM}} = \exp[-\tilde{\mathbf{c}}(\mathbf{X}_i)^T \hat{\lambda}]$ for all $i \in \{i: S_i = 1\}$. We can then use a Horvitz-Thompson type estimator similar to the inverse odds of sampling weights to estimate τ_{TATE} ,

$$\hat{\tau}_{\text{MOM}} = \sum_{\{i: S_i = 1\}} \frac{\hat{\gamma}_i^{\text{MOM}} (2Z_i - 1) Y_i}{\sum_{\{i: S_i = 1\}} \hat{\gamma}_i^{\text{MOM}} Z_i}.$$

In Signorovitch *et al.*, Assumptions 1–3 along with Assumption 5, are necessary to establish the consistency of $\hat{\tau}_{\text{MOM}}$ for τ_{TATE} .² More recent work can be adapted to show that this estimator is also consistent when Assumption 4 holds, thus achieving the doubly-robust property.²⁵

Our proposed adaptation to this entropy balancing solution instead sets $\mathbf{A}_i = (\mathbf{A}_{i0}^T, \mathbf{A}_{i1}^T)^T$ with $\mathbf{A}_{i0} = (1 - Z_i)\mathbf{c}(\mathbf{X}_i)$ and $\mathbf{A}_{i1} = Z_i\mathbf{c}(\mathbf{X}_i)$ while $\mathbf{B} = (\hat{\theta}_0^T, \hat{\theta}_1^T)^T$ to solve (6) using the following separable Lagrangian dual problem,

$$\begin{aligned} \hat{\lambda}_0 &= \arg \max_{\lambda \in \mathfrak{R}^m} \sum_{\{i: S_i = 1\}} \left\{ -\exp[-(1 - Z_i)\mathbf{c}(\mathbf{X}_i)^T \lambda] - \hat{\theta}_0^T \lambda \right\} \quad \text{and} \\ \hat{\lambda}_1 &= \arg \max_{\lambda \in \mathfrak{R}^m} \sum_{\{i: S_i = 1\}} \left[-\exp(-Z_i\mathbf{c}(\mathbf{X}_i)^T \lambda) - \hat{\theta}_1^T \lambda \right]. \end{aligned} \tag{8}$$

The empirical sampling weights are subsequently found with

$$\hat{\gamma}_i^{\text{CAL}} = \exp\left[-(1 - Z_i)\mathbf{c}(\mathbf{X}_i)^T \hat{\lambda}_0 - Z_i\mathbf{c}(\mathbf{X}_i)^T \hat{\lambda}_1\right] \text{ for all } i \in \{i: S_i = 1\}. \tag{9}$$

The estimator for τ_{TATE} using these estimated sampling weights is the same Horvitz-Thompson type estimator used by both the MOM and the IOSW approaches,

$$\hat{\tau}_{\text{CAL}} = \sum_{\{i: S_i = 1\}} \frac{\hat{\gamma}_i^{\text{CAL}}(2Z_i - 1)Y_i}{\sum_{\{i: S_i = 1\}} \hat{\gamma}_i^{\text{CAL}}Z_i}. \tag{10}$$

Notice that the covariate distributions are balanced between treatment groups and between the target sample and the trial participants. This is in contrast with the MOM estimator which only balances the covariate distribution between the target sample and the trial participants. This alteration to $\hat{\tau}_{\text{MOM}}$ remains doubly-robust for estimating τ_{TATE} given either Assumption 4 or 5. The double-robustness property of $\hat{\tau}_{\text{CAL}}$ is examined more closely in the Supplementary Material. The alteration to the MOM estimator is also motivated by the equivalence of (8)–(10) to the exponential tilting estimator implicitly proposed by Chan *et al.*, which is why we refer to it as the calibration (CAL) estimator.¹³ Recall that τ_{TATE} is a special case of a weighted average treatment. According to Theorem 3 of Chan *et al.*, if we can uniformly approximate $\rho(\mathbf{X}_i)$, $\mu_1(\mathbf{X}_i)$, and $\mu_0(\mathbf{X}_i)$ using a sufficiently rich basis represented by the balance functions $c_1(\mathbf{X}_i)$, $c_2(\mathbf{X}_i)$, ..., $c_m(\mathbf{X}_i)$ (i.e. the number of balancing functions m increases with n) while assuming mild conditions about the data generating processes, then the estimate of τ_{TATE} with $\hat{\tau}_{\text{CAL}}$ will achieve the efficiency bound in (1).¹³

This efficiency property is not shared by the method of moments version of entropy balancing, further motivating the calibration version.²⁵ Additional details on how to estimate the variance appears in the Supplementary Material.

There are a few reasons why we use the relative entropy over other criterion distance functions for transporting experimental results. The first is due to the resemblance of (9) to the inverse odds of sampling prescribed under Assumption 5. This relationship has been noted in several other articles.^{2,23} Another reason for using the relative entropy is the

guarantee that the estimated sampling weights will always be positive. Another suggestion might be to construct a Lagrangian dual using the Euclidean distance as the criterion function to get a different version of (6) and (8). However, the support for the Euclidean distance is the real numbers, implying that negative weights are feasible in such a setup. Adding the necessary constraint that $\gamma_i > 0$ for all $i = 1, 2, \dots, n$ makes the optimization problems in (6) and (8) less straightforward to solve.

Now consider the setting where we are provided only the sample covariate moments from the target sample. Assuming that $\hat{\theta}_0$ is fixed results in an inflated Type I error rate for inferences of τ_{TATE} . The one exception to this rule is when $\hat{\theta}_0 = \theta_0$ with zero variability. In other words, we would need to estimate $\hat{\theta}_0$ over the entire target population. If we are provided individual-level covariate data from the target sample, then we may derive a variance estimator for estimates of τ_{TATE} as opposed to τ_{TSATE} . Despite this shortfall, the estimators (8) – (10) remain consistent for τ_{TATE} in either setting. The same rule applies for both the OM approach and the MOM estimator since neither of these methods necessarily require the complete individual-level covariate data. A more concrete demonstration of this phenomenon is shown in Section 4.2.

4 | NUMERICAL EXAMPLES

4.1 | Simulation Study

In this section we present a simulation study to better understand the performance of entropy balancing techniques compared with the alternative methods illustrated in Section 2.3. We consider four experimental scenarios that test the consistency and efficiency of the estimators on finite-samples by altering the data generating processes. These scenarios all make the conditional constancy of absolute effects assumption defined in (2) to ensure compatibility between all the methods in consideration.

The first scenario establishes a baseline. For $i = 1, 2, \dots, n$, let $(X_{i0}|S_i = 0) \sim \mathcal{N}(-1, 4)$, $(X_{i1}|S_i = 0) \sim \text{Bin}(1, 0.6)$, $(X_{i2}|S_i = 0) \sim \mathcal{N}(0, 1)$, and $(X_{i3}|S_i = 0) \sim \text{Bin}(1, 0.5)$. Let $(X_{i0}|S_i = 1) \sim \mathcal{N}(1, 4)$, $(X_{i1}|S_i = 1) \sim \text{Bin}(1, 0.4)$, $(X_{i2}|S_i = 1) \sim \mathcal{N}(0, 1)$, and $(X_{i3}|S_i = 1) \sim \text{Bin}(1, 0.5)$. We generate the treatment assignment by sampling $Z_i \sim \text{Bin}(1, 0.5)$. The conditional mean of the potential outcomes are constructed as

$$\begin{aligned}\mu_0(\mathbf{X}_i) &= 10 - 3X_{i0} - X_{i1} + X_{i2} + 3X_{i3} \text{ and} \\ \mu_1(\mathbf{X}_i) &= \mu_0(\mathbf{X}_i) + 5 + 3X_{i0} - X_{i1} + X_{i2} - 3X_{i3}.\end{aligned}\tag{11}$$

Gaussian potential outcomes for each experimental scenario are generated by sampling $Y_{i(0)} \sim \mathcal{N}[\mu_0(\mathbf{X}_i), \sigma^2]$ and $Y_{i(1)} \sim \mathcal{N}[\mu_1(\mathbf{X}_i), \sigma^2]$, with the observed outcome determined by $Y_i = Z_i Y_{i(1)} + (1 - Z_i) Y_{i(0)}$ for each $i = 1, 2, \dots, n$. We discard the n_0 values of Y_i and Z_i for all $i \in \{i: S_i = 0\}$. We will refer to this set of conditions with the label “baseline”. Unless stated

otherwise, $S_i, \mathbf{X}_i = (X_{i0}, X_{i1}, X_{i2}, X_{i3}), Y_i$ and Z_i are provided to estimate $\hat{\gamma}_i^{\text{PS}}, \hat{\gamma}_i^{\text{MOM}}, \hat{\gamma}_i^{\text{CAL}}, \hat{\mu}_1(\mathbf{X}_i), \hat{\mu}_0(\mathbf{X}_i), \tilde{\mu}_1(\mathbf{X}_i),$ and $\tilde{\mu}_0(\mathbf{X}_i)$ which are required in the estimators described in 2.3 and 3.

In the scenarios labeled “positivity”, we increase the difference between the two covariate distributions by substituting $(X_{i0}|S_i = 0) \sim \mathcal{N}(1, 2), (X_{i0}|S_i = 1) \sim \mathcal{N}(-1, 1), (X_{i1}|S_i = 0) \sim \text{Bin}(1, 0.2),$ and $(X_{i1}|S_i = 1) \sim \text{Bin}(1, 0.8)$ for the respective covariates into the data generating mechanisms. This alteration will test the sensitivity of each method on the intrinsic limitations associated with Assumption 2. For the scenario labeled “sparse”, we provide each method an additional set of covariates that do not affect the responses. The potential outcomes are still determined from (11) with the original covariate values, yet the different estimators must also accommodate the additional noise covariates of $(X_{i,r}|S_i = 0) \sim (X_{i,r-4})|S_i = 1$ and $(X_{i,r}|S_i = 1) \sim (X_{i,r-4})|S_i = 0$ for $r \in \{4, 5, 6, 7\}$. Each of the estimators are provided data for $(X_{i0}, X_{i1}, \dots, X_{i7})$ in addition to Y_i, Z_i and S_i . In the scenarios labeled “missing”, we generate the outcomes according to (11) yet we provide each method only (X_{i0}, X_{i2}) and omit X_{i1} and X_{i3} while estimating τ_{TATE} . Note that this means we omit one of the effect modifiers, X_{i1} . Next, we formulate scenarios which misspecify the outcome model (“outcome”). To do so, we generate outcomes according to the model

$$\begin{aligned} \mu_0(\mathbf{X}_i) &= 10 - 3U_{i0} - X_{i1} + U_{i2} + 3X_{i3} \text{ and} \\ \mu_1(\mathbf{X}_i) &= \mu_0(\mathbf{X}_i) + 5 + 3U_{i0} - X_{i1} + U_{i2} - 3X_{i3} \end{aligned} \tag{12}$$

where $U_{i0} = \exp(-X_{i0}/4 + X_{i2}/4)$ and $U_{i2} = (X_{i0}/2 - X_{i2}/2)^2$. Both U_{i0} and U_{i2} are standardized across both samples to have a mean of 0 and variance 1 in the combined trial and target samples. We then provide each method the original covariate values $(X_{i0}, X_{i1}, X_{i2}, X_{i3})$ for all $i = 1, 2, \dots, n$. On the other hand, in the sampling misspecification scenario (“sampling”), we provide each method data for $(U_{i0}, X_{i1}, U_{i2}, X_{i3})$ while the outcomes are still generated by the model in (12). The standardization step is key to ensure that the true magnitude of the differences between the sample covariate distributions are never fully expressed by the sampling model. In addition to varying the scenarios that test the violations to the assumptions in Section 2.2, we also vary $n_0 \in \{500, 1000\}$ and $n_1 \in \{500, 1000\}$, creating 24 different conditions for which we will generate 1000 replications.

We report the average bias and empirical mean squared error of the average treatment effect estimates across the 1000 iterations for each scenario. The average model and empirical standard errors are provided in tables S1 and S2 in the Supplementary Materials. The model standard errors are obtained using a sandwich variance estimator for each estimate in every iteration of the simulation. The empirical standard errors are the standard deviations of the estimates from each estimator pooled across the iterations of a given scenario. The five methods we compare for estimating the target average treatment effect are: Inverse Odds of Sampling Weights (IOSW), G-Computation (OM), Augmented Inverse Odds of Sampling Weights (DR), Targeted Maximum Likelihood Estimation (TMLE), Method of Moments (MOM), and Calibration (CAL). Additionally, for IOSW, OM, DR, and TMLE, both $\hat{\mu}_1(\mathbf{X}_i)$ and $\hat{\mu}_0(\mathbf{X}_i)$ are fit by regressing Y_i onto the covariates provided in each scenario with data from $S_i = 1$ and stratified by the Z_i . $\hat{\rho}(\mathbf{X}_i)$ is fit with logistic regression using covariates that

predict S_j . The standard errors are estimated using robust sandwich variance estimators. For the non-entropy balancing type estimators, R code for finding variance estimates are provided in the existing literature.^{1,16}

The average bias and mean squared errors of the experiment are summarized in Table 1. A visual comparison for a subset of the results featured in Table 1 where $n_1 = 200$ and $n_0 = 500$ appear in the boxplots of Figure 2. Each method produces consistent estimates under the baseline scenario. However, each method also has its short-comings. First, we can see that IOSW produce highly variable estimates in cases where the positivity assumption (Assumption 2) is practically violated and is biased in cases when the sampling scenario is misspecified. On the other hand, the OM approach is biased when the outcome model is misspecified. TMLE, DR, MOM, and CAL all appear to produce unbiased estimates of the average treatment effect in every scenario. This is interesting for the MOM estimator since this would imply that it is also doubly-robust in terms of consistency. Some insight into why this might be is provided elsewhere.²⁵ However, we can see in Table 1 that CAL had either the same or smaller mean squared errors over TMLE and MOM. In some scenarios DR did have smaller errors than CAL. Nevertheless, The OM approach had the smallest mean squared errors across most scenarios, other than in the scenarios where we misspecify the outcome model naturally. When the models miss (or ignore) some of the effect modifiers, we see that every method we test produces biased estimates of the target population average treatment effect. This particular scenario emphasizes the results of these estimators when both the outcome and the sampling models are misspecified, even when missing a single effect modifier.

There is a downside to the so-called calibration version of entropy balancing. In the sparse and positivity violation scenarios, the number of models that converge decreases considerably. When $n_1 = 200$ and $n_0 = 500$, CAL was only able to find a solution in 64.2% of the iterations under the sparse scenario and 49.0% of the iterations when positivity is practically violated. When $n_1 = 200$ and $n_0 = 1000$ we observe a 66.0% and 48.1% rate of convergence in the sparse and positivity scenarios, respectively. Otherwise, the calibration approach to entropy balancing converged in each iteration for every other scenario. Meanwhile, the method of moments approach to entropy balancing also failed to converge in approximately 7.5% of the scenarios in both of the positivity scenarios where the calibration estimator often failed to converge.

4.2 | Coverage Probabilities of τ_{TATE} and τ_{TSATE}

Consider the baseline scenario in the previous set of simulations. Using the individual-level data from the trial sample, and the target sample covariate moments, we demonstrate how inferences for τ_{TATE} will have an inflated Type I error. We do so by finding the empirical coverage probability of both τ_{TSATE} and τ_{TATE} with both of the entropy balancing approaches described in Section 3. Robust sandwich variance estimators are used to construct the confidence intervals. The coverage probability is obtained by averaging over the indicator variable generated by whether the resulting confidence interval about the average treatment effect estimate covers either of τ_{TSATE} or τ_{TATE} at each iteration. This will demonstrate why entropy balancing can only be used to infer upon the target sample average

treatment effect instead of the target population average treatment effect unless the entire individual-level data about the balance functions in both the target and trial samples is available. For this simulation experiment, we set the target and trial sample sizes at $n_0 \in \{500, 1000, 10000\}$ and $n_1 \in \{1000, 10000\}$, respectively. We use large sample sizes to ensure the accuracy of the robust variance estimator.

The results in Table 2 show how modifying n_1 and n_0 affects the coverage probabilities for τ_{TSATE} and τ_{TATE} for the setting where we are given the target sample covariate moments. Observe that the coverage probability of τ_{TSATE} is dependent on n_1 alone - as n_1 increases, the coverage probabilities increase. The coverage probability of τ_{TATE} , on the other hand, appears to be dependent on the ratio between n_0 and n_1 . For inference on τ_{TATE} , we see the best results occur when n_1 is small relative to n_0 . When $n_1 = 1000$ and $n_0 = 10000$, the variation of $\hat{\theta}$ has less impact on the total variance, producing the best results. In contrast, when $n_1 = 10000$ and $n_0 = 10000$, the variation of $\hat{\theta}$ has a greater impact, resulting in a decreased probability of coverage. This observation is only compounded in cases where $n_1 > n_0$. This leads us to believe that n_1 needs to be sufficiently large while also remaining small compared to n_0 in order to be effective for inferring on τ_{TATE} . When we adjust the sandwich estimator to incorporate individual-level covariate data from the target sample, we see that the accuracy of coverage probability is now tied to the total sample size $n = n_0 + n_1$, which is typical for robust variance estimators as they are derived under asymptotic conditions.

5 | TRANSPORTING RESULTS OF ACCORD-BP STUDY TO THE US POPULATION

Translating clinical trial results to clinical care is particularly challenging when the results of two studies conducted for similar indications and treatments yield conflicting conclusions. For example, the optimal approach to hypertension treatment remains unclear, partly due to conflicting clinical trial results. The Systolic Blood Pressure Intervention Trial (SPRINT) and the Action to Control Cardiovascular Risk in Diabetes Blood Pressure (ACCORD-BP) trial both randomized participants with hypertension to intensive (< 120 mmHg) or conventional (< 140 mmHg) blood pressure control targets. The study populations differed in that ACCORD-BP was limited to diabetes patients while SPRINT excluded diabetes patients. The two similarly designed studies in differing populations had different results: SPRINT, but not ACCORD-BP, found an association of intensive blood pressure control with several clinically meaningful outcomes including cardiovascular disease events.^{26,27} Importantly, the ACCORD-BP trial was enriched for individuals at high cardiovascular disease risk aside from the presence of diabetes, raising the question of whether the result of the trial applies to a more general diabetes patient population. Thus, transporting the ACCORD-BP trial to the general US population of diabetes patients may provide insight into hypertension management for individuals with diabetes and help reconcile the discrepant trial results.

To address this question, Berkowitz *et al.* used inverse odds of sampling weights (IOSW) to transport the ACCORD-BP trial to a sample of US diabetes patients drawn from the US National Health and Nutrition Examination Survey (NHANES).²⁸ They found that

weighting the ACCORD-BP sample to reflect the diabetes patient sample in NHANES yielded intervention effects more similar to those observed in the SPRINT trial of non-diabetes patients than in the unweighted ACCORD-BP trial. We use this previously demonstrated application of transportability methods to ACCORD-BP as a real-world application of the entropy balancing (CAL) methods described here. In our applied example, we transport four-year post randomization risk difference estimates of total mortality observed in the ACCORD-BP trial²⁶ to a sample of US diabetes patients drawn from the NHANES cohort.²⁸ We use two methods for transporting the results of ACCORD-BP to NHANES - IOSW and entropy balancing (CAL). Furthermore, using entropy balancing we provide confidence intervals about the target sample average treatment effect and the target population average treatment effect. Recall that the former estimand does not require any individual-level data from the NHANES sample.

Table 3 shows covariates (which we set as the balance functions) balanced between ACCORD-BP and NHANES, their unweighted sample covariate moments from the NHANES and from the ACCORD-BP data, and the subsequent weighted covariate sample moments of the ACCORD-BP sample after balancing. Compared to ACCORD-BP, the NHANES diabetes sample was younger, more likely to be Hispanic and less likely to be black, more likely to be never smokers, less likely to have a history of myocardial infarction (MI) but more likely to have a history of congestive heart failure (CHF), and had a shorter duration of diabetes and better glycemic control (indicated by hemoglobin A1c) (Table 3). Many of the differences in covariate distributions reflect that ACCORD trial eligibility criteria focused on those with relatively long duration of diabetes and high prevalence of cardiovascular risk factors. Of note, the intensive blood pressure control intervention had a smaller benefit in individuals with pre-existing cardiovascular disease in the SPRINT trial, making it plausible that differences between the ACCORD-BP population and a general population of diabetes patients might modify the effect of the blood pressure intervention.²⁷ In another study using data from NHANES, hemoglobin A1c was associated with increased risk of all-cause and cause-specific mortality.²⁹ Zoungas *et al.*³⁰ show that diabetes duration is associated with death while McEwen *et al.*³¹ identified multiple predictors of total mortality such as race, age, and previous cardiovascular events among diabetic patients. These previous findings imply that numerous factors might have the potential for confounding the relationship between sampling and the outcome. The differences in baseline covariates between ACCORD-BP and NHANES are reduced after balancing with both CAL and IOSW. However, the covariate sample moments after CAL weighting consistently matched the NHANES sample more closely than after IOSW weighting (Table 3, Figure 3). Small residual differences remain between NHANES and the weighted ACCORD-BP sample, for example with triglycerides and high density lipoproteins (Figure 3).

The ACCORD-BP study originally found an increase in four-year mortality of 0.59% [95% CI:(-0.75%, 1.93%)] in the intensive treatment group. After weighting the ACCORD-BP responses with inverse odds of sampling weights estimated with maximum likelihood, the estimated risk difference on the NHANES population is -1.35% [95% CI: (-3.5%, 0.8%)]. Using CAL, we observe a risk difference of -0.04% [95% CI: (-1.80%, 1.71%)] where the confidence interval corresponds to the NHANES sample average treatment effect. The 95%

confidence interval for the NHANES population average treatment effect is (−1.94%, 1.86%) when using the individual level covariate data from the NHANES sample.

Though the total mortality is insignificant at a 0.05 level of significance, regardless of method, we see changes in the risk difference estimate. The original analysis found an increase in mortality among the intensively treated patients. IOSW weights yielded a decreased total mortality among intensively treated patients in the NHANES population, while CAL weights yielded a nearly null result. These differences seem to indicate the presence of effect modifiers contributing to the effect of blood pressure treatment intensity on mortality. Notice also that the population level estimate is the same as the sample level estimate using entropy balancing. However, the width of the confidence interval is wider for the population level estimate. Nevertheless, the population level estimate from the CAL approach is still narrower than the estimated confidence interval produced by the IOSW approach, indicating an increase in efficiency.

6 | DISCUSSION

In this article we have described a doubly-robust method for transporting experimental results borrowed from the entropy optimization literature. We also borrow results from the indirect comparison literature, which allows us to relax the conditional constancy of absolute effects assumption typically applied in the transportability literature and focus our efforts on modeling relative effects with effect modifiers rather than the absolute effects which require both the effect modifiers and any prognostic variables.^{1,16} However, if the sampling model is incorrect, then we would need the conditional constancy of absolute effects assumption to hold in order to get consistent estimates given the doubly-robust property. As a result more emphasis should be placed on correctly specifying the sampling model over the outcome model if modelling choices begin to deviate. The entropy balancing methodology may operate in two settings - when we are presented with the complete individual-level data of the trial sample and either the individual-level covariate data or the covariate sample moments of the target sample. The distinction between the two settings amounts to inferring upon the target population average treatment effect versus the target sample average treatment effect. We showed entropy balancing to be an efficient causal effect estimator in finite-samples through simulation. We also compared two methods for transporting the ACCORD-BP study to the NHANES population. These numerical examples demonstrate some of the practical implications of our work.

The drawback to using entropy balancing for transportability is with the algorithm's rate of convergence. In small samples, the probability that a feasible weighting solution exists decreases, particularly when the positivity assumption is practically violated. One solution applied to covariate balance problems is to use inequality constraints to mitigate treatment group heterogeneity.⁹ This solution is most useful in high-dimensional settings. There may also be a way to incorporate the method of moment balancing weights into the TMLE framework by substituting $\hat{\gamma}^{\text{MOM}}$ for $\hat{\gamma}^{\text{PS}}$ in (5). This could eventually set up a targeted maximum likelihood-type estimator that can operate in the setting where we do not have any individual-level data from the target population. In the case where convergence failure occurs due to a high dimension of potential effect modifiers relative to the trial sample size,

one should carefully consider balance diagnostics between trial participants and non-participants in much the same way as one identifies potential confounders in an observational study.^{32,33} Additional sensitivity analyses should be employed intermittently to ensure that every effect modifier is accounted for.³⁴ In the case of positivity violations, it seems that methods that employ more direct implementations of an outcome model, such as the G-computation and DR approaches, fair better given their ability to extrapolate over the covariate space. Violations of Assumptions 1 and 3 pose a more difficult challenge to evaluate as these assumptions are untestable. Expert level knowledge of the domain area are necessary to ensure that these assumptions will hold with the preferred transportability model.

Future work will address two additional data settings not evaluated here. First, the setting where the target sample contains data from a second randomized experiment, including both the individual-level outcome and the treatment assignment. The process of combining experiments, termed as data-fusion, is beyond what we discuss in this paper but is nevertheless an important problem which we would like to approach with entropy balancing in future research. A second direction for future work is to examine methods for transportability between two observational samples, rather than assuming availability of randomized clinical trial data for the trial sample.³⁵ In this situation, we would also need to model the probability of treatment within the the observational study representing the “trial” sample. We might also seek to relax Assumptions 4 and 5 using a nonparametric setup to the problem similar to the sieve approach but instead applied to transportability.^{12,13} Finally, while the average treatment effect estimands under consideration in this manuscript are applicable to various outcomes, including a binary one, more work is needed to generalize many of these estimators to accommodate a non-linear link function for the outcome model.

In summary, entropy balancing provides an approach to transportability that is flexible regarding the applicable data settings and exhibits double robustness in specific scenarios. In particular, entropy balancing yields more precise effect estimates across a range of simulation scenarios when the target population is large than alternative methods using only covariate sample moments from the target population.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

Funding information:

This research was supported in part by the US Department of Veterans Affairs Award IK2-CX001907-01. S.A. Berkowitz’s role in the research reported in this publication was supported by the National Institute Of Diabetes And Digestive And Kidney Diseases of the National Institutes of Health under Award Number K23DK109200. D. Ghosh’s role reported in the publication was supported by NSF DMS-1914937.

APPENDIX

A SIMULATION CODE

Code for reproducing the simulation experiment conducted in Section 4.1 is available at the following address: <https://github.com/kevjosey/transport-sim>. All data analyzed in this study are publicly available to investigators with approved human subjects approval via the US National Institutes of Health, National Heart, Lung, and Blood Institutes, Biologic Specimen and Data Repository Information Coordinating Center (ACCORD Study, <https://biolincc.nhlbi.nih.gov/studies/accord/>) or the US National Center for Health Statistics (NHANES, <https://www.cdc.gov/nchs/nhanes/>). Statistical code for creating analytic datasets and for performing analyses are available from the authors upon request.

References

1. Rudolph KE, van der Laan MJ. Robust estimation of encouragement design intervention effects transported across sites. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*2017; 79(5): 1509–1525. 10.1111/rssb.12213 [PubMed: 29375249]
2. Signorovitch JE, Wu EQ, Yu AP, et al. Comparative Effectiveness Without Head-to-Head Trials. *PharmacoEconomics*2010; 28(10): 935–945. 10.2165/11538370-000000000-00000 [PubMed: 20831302]
3. Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. *Biometrics*2005; 61(4): 962–973. 10.1111/j.1541-0420.2005.00377.x [PubMed: 16401269]
4. Hartman E, Grieve R, Ramsahai R, Sekhon JS. From sample average treatment effect to population average treatment effect on the treated: combining experimental with observational studies to estimate population treatment effects. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*2015; 178(3): 757–778. 10.1111/rssa.12094
5. Zhang Z, Nie L, Soon G, Hu Z. New methods for treatment effect calibration, with applications to non-inferiority trials. *Biometrics*2016; 72(1): 20–29. 10.1111/biom.12388 [PubMed: 26363775]
6. Phillippo DM, Ades AE, Dias S, Palmer S, Abrams KR, Welton NJ. Methods for Population-Adjusted Indirect Comparisons in Health Technology Appraisal. *Medical Decision Making*2018; 38(2): 200–211. 10.1177/0272989X17725740 [PubMed: 28823204]
7. Hainmueller J. Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies. *Political Analysis*2012; 20(1): 25–46. 10.1093/pan/mpr025
8. Imai K, Ratkovic M. Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*2014; 76(1): 243–263. 10.1111/rssb.12027
9. Wang Y, Zubizarreta JR. Minimal dispersion approximately balancing weights: asymptotic properties and practical considerations. *Biometrika*2020; 107(1): 93–105. 10.1093/biomet/asz050
10. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*1974; 66(5): 688–701. 10.1037/h0037350
11. Imai K, King G, Stuart EA. Misunderstandings between experimentalists and observationalists about causal inference. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*2008; 171(2). 10.1111/j.1467-985X.2007.00527.x
12. Hirano K, Imbens GW, Ridder G. Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score. *Econometrica*2003; 71(4): 1161–1189. 10.1111/1468-0262.00442
13. Chan KCG, Yam SCP, Zhang Z. Globally efficient non-parametric inference of average treatment effects by empirical balancing calibration weighting. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*2016; 78(3): 673–700. 10.1111/rssb.12129 [PubMed: 27346982]
14. Deville JC, Särndal CE. Calibration Estimators in Survey Sampling. *Journal of the American Statistical Association*1992; 87(418): 376–382. 10.1080/01621459.1992.10475217

15. Pearl J, Bareinboim E. External Validity: From Do-Calculus to Transportability Across Populations. *Statistical Science*2014; 29(4): 579–595.
16. Dahabreh IJ, Robertson SE, Steingrimsson JA, Stuart EA, Hernán MA. Extending inferences from a randomized trial to a new target population. *Statistics in Medicine*2020; 39(14): 1999–2014. 10.1002/sim.8426 [PubMed: 32253789]
17. Rosenbaum PR, Rubin DB. The central role of the propensity score in observational studies for causal effects. *Biometrika*1983; 70(1): 41–55. 10.1093/biomet/70.1.41
18. Phillippo DM, Dias S, Ades AE, Welton NJ. Equivalence of entropy balancing and the method of moments for matching-adjusted indirect comparison. *Research Synthesis Methods*2020; 11(4): 568–572. [PubMed: 32395870]
19. Westreich D, Edwards JK, Lesko CR, Stuart E, Cole SR. Transportability of Trial Results Using Inverse Odds of Sampling Weights. *American Journal of Epidemiology*2017; 186(8): 1010–1014. 10.1093/aje/kwx164 [PubMed: 28535275]
20. Robins JA new approach to causal inference in mortality studies with a sustained exposure period —application to control of the healthy worker survivor effect. *Mathematical Modelling*1986; 7(9): 1393–1512. 10.1016/0270-0255(86)90088-6
21. Robins JM, Rotnitzky A, Zhao LP. Estimation of Regression Coefficients When Some Regressors are not Always Observed. *Journal of the American Statistical Association*1994; 89(427): 846–866. 10.1080/01621459.1994.10476818
22. van der Laan MJ, Rubin D. Targeted Maximum Likelihood Learning. *The International Journal of Biostatistics*2006; 2(1). 10.2202/1557-4679.1043
23. Zhao Q, Percival D. Entropy Balancing is Doubly Robust. *Journal of Causal Inference*2017; 5(1). 10.1515/jci-2016-0010
24. Josey KP, Juarez-Colunga E, Yang F, Ghosh D. A Framework for Covariate Balance using Bregman Distances. *Scandinavian Journal of Statistics*2020. 10.1111/sjos.12457
25. Dong L, Yang S, Wang X, Zeng D, Cai J. Integrative analysis of randomized clinical trials with real world evidence studies. 2020. <https://arxiv.org/abs/2003.01242v1>.
26. ACCORD Study Group . Effects of Intensive Blood-Pressure Control in Type 2 Diabetes Mellitus. *New England Journal of Medicine* 2010; 362(17): 1575–1585. 10.1056/NEJMoa1001286
27. SPRINT Research Group. A Randomized Trial of Intensive versus Standard Blood-Pressure Control. *New England Journal of Medicine*2015; 373(22): 2103–2116. 10.1056/NEJMoa1511939
28. Berkowitz SA, Sussman JB, Jonas DE, Basu S. Generalizing Intensive Blood Pressure Treatment to Adults With Diabetes Mellitus. *Journal of the American College of Cardiology*2018; 72(11): 1214–1223. 10.1016/j.jacc.2018.07.012 [PubMed: 30189998]
29. Palta P, Huang ES, Kalyani RR, Golden SH, Yeh HC. Hemoglobin A1c and Mortality in Older Adults With and Without Diabetes: Results From the National Health and Nutrition Examination Surveys (1988–2011). *Diabetes Care*2017; 40(4): 453–460. 10.2337/dci16-0042 [PubMed: 28223299]
30. Zoungas S, Woodward M, Li Q, et al. Impact of age, age at diagnosis and duration of diabetes on the risk of macrovascular and microvascular complications and death in type 2 diabetes. *Diabetologia*2014; 57(12): 2465–2474. 10.1007/s00125-014-3369-7 [PubMed: 25226881]
31. McEwen LN, Karter AJ, Waitzfelder BE, et al. Predictors of mortality over 8 years in type 2 diabetic patients: Translating Research Into Action for Diabetes (TRIAD). *Diabetes Care*2012; 35(6): 1301–1309. 10.2337/dc11-2281 [PubMed: 22432119]
32. Brookhart MA, Schneeweiss S, Rothman KJ, Glynn RJ, Avorn J, Stürmer T. Variable Selection for Propensity Score Models. *American Journal of Epidemiology*2006; 163(12): 1149–1156. 10.1093/aje/kwj149 [PubMed: 16624967]
33. Westreich D, Cole SR, Funk MJ, Brookhart MA, Stürmer T. The role of the c-statistic in variable selection for propensity score models. *Pharmacoepidemiology and Drug Safety*2011; 20(3): 317–320. 10.1002/pds.2074 [PubMed: 21351315]
34. Nguyen TQ, Ackerman B, Schmid I, Cole SR, Stuart EA. Sensitivity analyses for effect modifiers not observed in the target population when generalizing treatment effects from a randomized controlled trial: Assumptions, models, effect scales, data scenarios, and implementation details. *PLOS ONE*2018; 13(12): e0208795. 10.1371/journal.pone.0208795 [PubMed: 30533053]

35. Josey KP, Yang F, Ghosh D, Raghavan S. A Calibration Approach to Transportability with Observational Data. arXiv preprint arXiv:2008.066152020.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

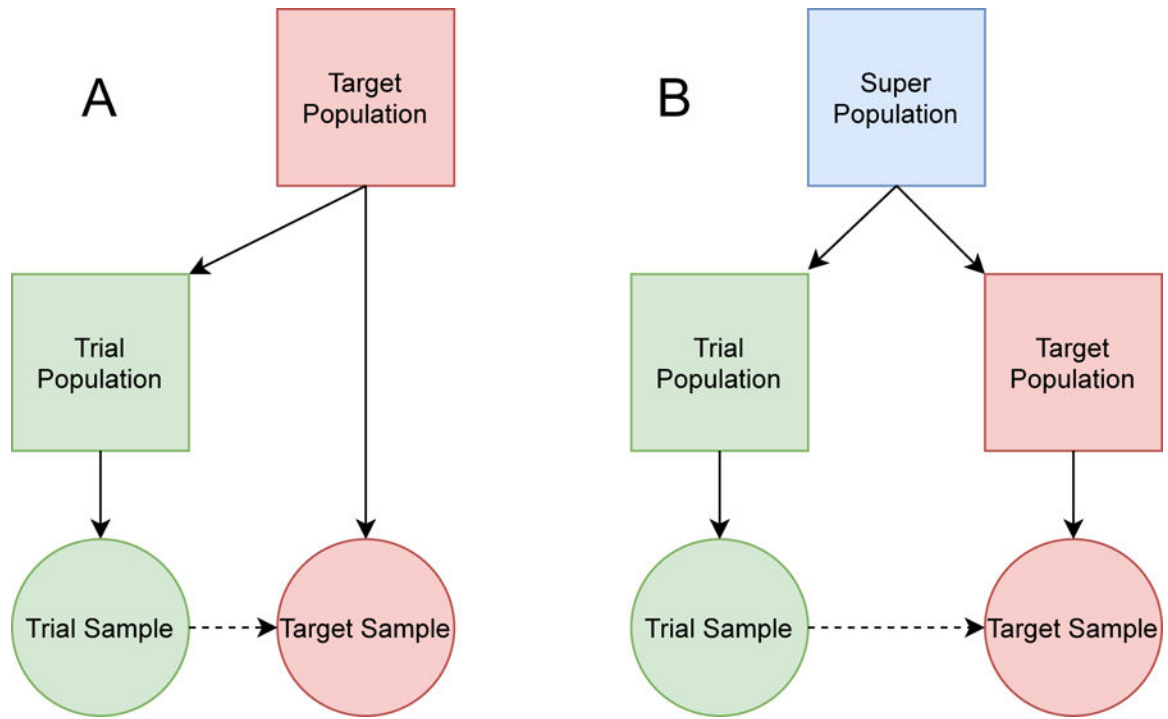


FIGURE 1. Square nodes represent populations whereas circular nodes represent samples. The solid arrow represents a subsetting of the origin node. The dashed line represents the process of generalizability (A) and transportability (B).

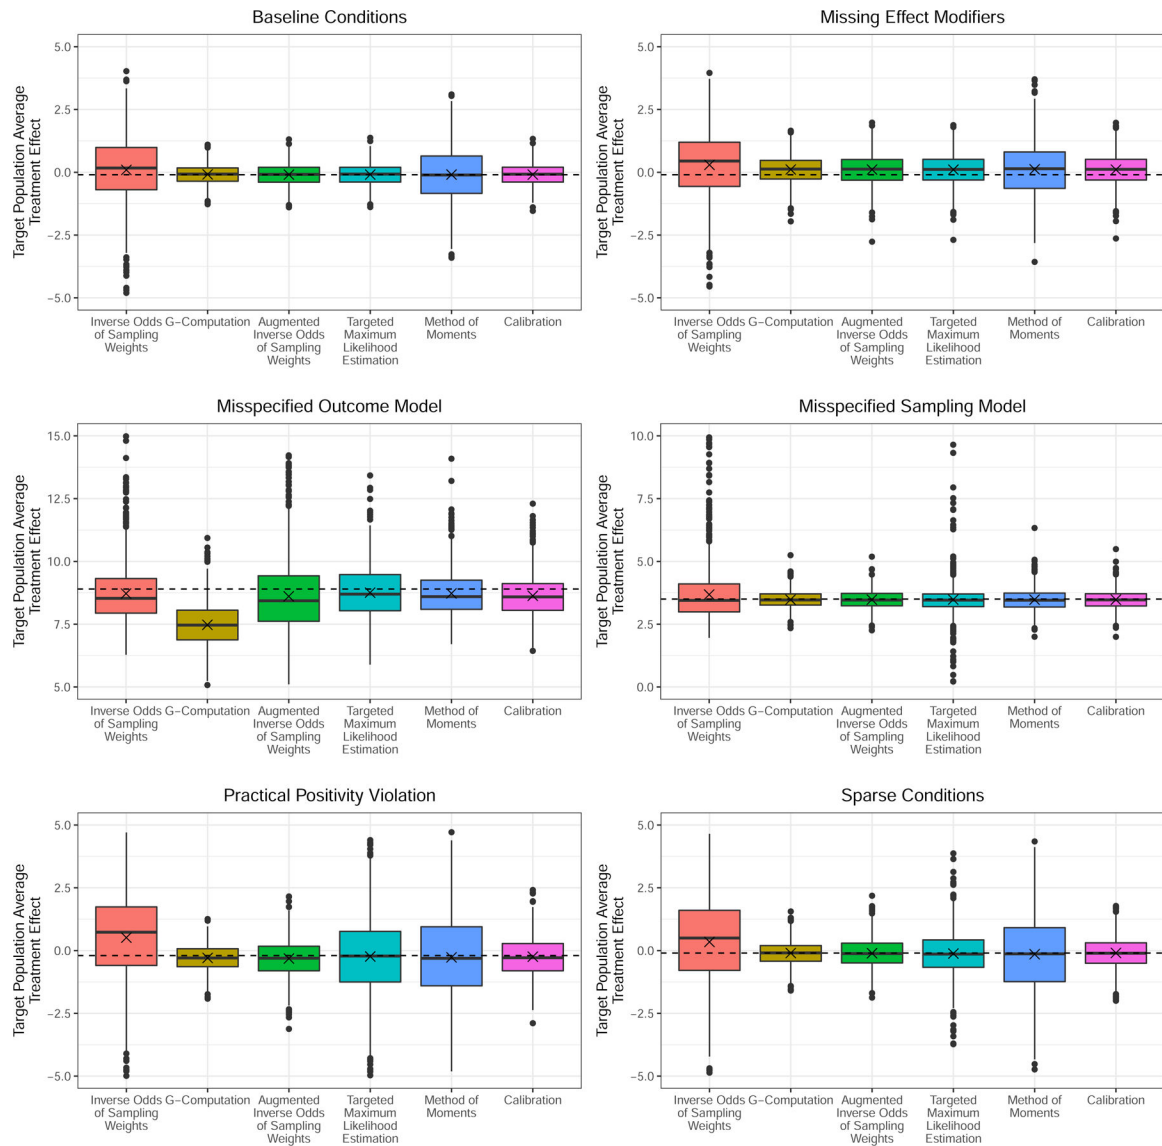


FIGURE 2. Estimates of the target population average treatment effects over the 1000 iterations of the simulation study described in Section 4.1. The dashed line demarcates the true target population average treatment effect for each scenario while the x is the average of the estimates. These estimates are drawn from cases when $n_1 = 200$ and $n_0 = 500$.




FIGURE 3. Absolute standardized mean differences for various weighting estimators between NHANES and ACCORD. The red dotted line demarcates an absolute standardized mean difference of 0.1.

TABLE 1

Average bias and (mean squared error) of τ_{TATE} estimates. The scenarios where $n_1 = 200$ and $n_0 = 500$ are detailed in Figure 2.

n_0	n_1	Scenario	τ_{TATE}	IOSW	OM	DR	TMLE	MOM	CAL
500	1000	baseline	-0.1	0.03 (0.47)	0.00 (0.09)	0.00 (0.10)	0.00 (0.10)	0.00 (0.33)	0.00 (0.10)
500	1000	missing	-0.1	0.23 (0.48)	0.19 (0.16)	0.19 (0.17)	0.19 (0.17)	0.19 (0.36)	0.19 (0.17)
500	1000	outcome	8.9	-0.03 (0.41)	-1.44 (2.23)	-0.08 (0.72)	-0.03 (0.26)	-0.05 (0.24)	-0.08 (0.20)
500	1000	positivity	-0.3	0.27 (1.51)	-0.01 (0.08)	-0.01 (0.19)	-0.01 (0.38)	0.00 (0.81)	-0.01 (0.19)
500	1000	sample	4.5	0.64 (3.61)	0.00 (0.03)	0.00 (0.05)	-0.02 (0.59)	0.00 (0.06)	0.00 (0.04)
500	1000	sparse	-0.1	0.08 (1.18)	-0.01 (0.11)	-0.01 (0.16)	-0.02 (0.20)	-0.03 (0.63)	-0.01 (0.16)
1000	200	baseline	-0.1	0.15 (2.04)	0.00 (0.14)	0.00 (0.18)	-0.01 (0.18)	-0.02 (1.18)	-0.01 (0.18)
1000	200	missing	-0.1	0.30 (1.94)	0.18 (0.27)	0.20 (0.33)	0.20 (0.33)	0.19 (1.07)	0.20 (0.33)
1000	200	outcome	8.9	-0.10 (1.66)	-1.45 (2.87)	-0.22 (2.87)	-0.15 (1.14)	-0.15 (0.87)	-0.26 (0.76)
1000	200	positivity	-0.3	0.91 (4.80)	0.03 (0.24)	0.05 (0.49)	0.05 (3.10)	0.11 (3.17)	0.08 (0.62)
1000	200	sample	3.2	0.37 (2.63)	-0.01 (0.10)	-0.01 (0.14)	0.07 (1.57)	0.00 (0.20)	-0.01 (0.13)
1000	200	sparse	-0.1	0.31 (4.14)	-0.01 (0.20)	-0.01 (0.37)	-0.01 (0.84)	-0.05 (2.67)	0.00 (0.41)
1000	1000	baseline	-0.1	0.01 (0.48)	0.00 (0.06)	0.00 (0.07)	0.00 (0.07)	-0.02 (0.29)	0.00 (0.07)
1000	1000	missing	-0.1	0.21 (0.46)	0.19 (0.11)	0.19 (0.13)	0.19 (0.13)	0.19 (0.30)	0.19 (0.13)
1000	1000	outcome	8.9	-0.03 (0.43)	-1.41 (2.14)	-0.06 (0.75)	-0.03 (0.23)	-0.05 (0.23)	-0.08 (0.20)
1000	1000	positivity	-0.3	0.19 (1.74)	0.00 (0.06)	0.00 (0.16)	-0.01 (0.31)	-0.01 (0.81)	0.00 (0.16)
1000	1000	sample	4.0	0.78 (4.22)	0.00 (0.02)	-0.01 (0.04)	-0.12 (1.89)	-0.01 (0.04)	-0.01 (0.03)
1000	1000	sparse	-0.1	0.12 (1.03)	-0.01 (0.07)	-0.01 (0.12)	-0.01 (0.14)	0.00 (0.57)	-0.01 (0.11)

TABLE 2

Coverage Probabilities of τ_{TSATE} and τ_{TATE} using Entropy Balancing and Outcome Modeling Techniques.

n_0	n_1	Without Individual Level Data				With Individual Level Data	
		MOM τ_{TSATE}	MOM τ_{TATE}	CAL τ_{TSATE}	CAL τ_{TATE}	MOM τ_{TATE}	CAL τ_{TATE}
500	1000	0.926	0.882	0.929	0.623	0.937	0.943
500	10000	0.956	0.667	0.936	0.289	0.948	0.960
1000	1000	0.922	0.897	0.938	0.761	0.928	0.951
1000	10000	0.929	0.770	0.938	0.372	0.945	0.942
10000	1000	0.928	0.923	0.930	0.902	0.923	0.929
10000	10000	0.930	0.912	0.960	0.791	0.935	0.958

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 3

Values are mean \pm SD or %. Means and percentages for NHANES are nationally representative using NHANES sampling weights.

Variables	NHANES	ACCORD-BP	IOSW ACCORD-BP	CAL ACCORD-BP
Baseline age, yrs	59.65 \pm 13.70	62.84 \pm 6.74	61.50 \pm 6.66	59.61 \pm 6.91
Female	48.9	47.1	48.0	48.9
Race/Ethnicity				
Non-Hispanic white	62.6	58.7	51.6	62.5
Non-Hispanic black	15.2	24.0	19.6	15.2
Hispanic	15.2	6.8	22.3	15.2
Asian/multi/other	7.0	10.5	6.5	7.1
Insurance	86.8	85.0	84.6	86.6
Smoking status				
Never	51.4	44.6	52.9	51.4
Former	33.1	42.6	30.9	33.1
Current	15.5	12.8	16.2	15.5
Education				
Less than HS	25.7	16.3	30.9	25.7
HS diploma	27.1	27.0	26.4	27.1
Some college	29.3	32.4	26.5	29.3
College diploma or higher	17.9	24.3	16.2	17.9
History of CHF	7.7	4.2	11.0	7.7
History of MI	10.5	13.6	11.4	10.5
History of stroke	7.9	6.4	7.5	7.8
Years with diabetes	7.49 \pm 9.20	10.88 \pm 7.83	10.05 \pm 7.26	7.50 \pm 6.51
BMI, kg/m ²	32.80 \pm 7.31	32.10 \pm 5.47	32.07 \pm 5.52	32.80 \pm 5.78
SBP, mm Hg	130.05 \pm 19.15	139.33 \pm 15.61	133.94 \pm 14.42	129.67 \pm 13.98
DBP, mm Hg	69.50 \pm 12.96	75.86 \pm 10.28	71.53 \pm 9.57	69.44 \pm 9.55
HDL, mg/dl	49.11 \pm 13.46	46.049 \pm 13.68	51.60 \pm 17.24	49.08 \pm 17.72
LDL, mg/dl	103.83 \pm 36.03	110.70 \pm 36.52	105.42 \pm 31.33	103.59 \pm 33.31
Triglycerides, mg/dl	148.93 \pm 76.13	193.36 \pm 174.21	125.40 \pm 68.01	147.21 \pm 95.52
FPG, mg/dl	151.88 \pm 54.62	174.81 \pm 57.66	171.54 \pm 57.47	151.11 \pm 47.30
HbA _{1c} , %	7.16 \pm 1.64	8.34 \pm 1.09	7.94 \pm 0.95	7.16 \pm 0.75
Estimated GFR, ml/min	87.46 \pm 28.11	91.64 \pm 29.83	84.99 \pm 21.16	87.31 \pm 22.66
Urine albumin to creatinine ratio	75.44 \pm 481.68	93.84 \pm 333.81	105.89 \pm 427.60	45.32 \pm 204.57