

# **rbw**: An R package for Constructing Residual Balancing Weights

Derick Baum, Harvard University      Xiang Zhou, Harvard University

September 10, 2021

## **Abstract**

We describe the R package **rbw**, which implements the method of residual balancing weights (RBW) for estimating marginal structural models. In contrast to other methods such as inverse probability weighting (IPW) and covariate balancing propensity scores (CBPS), RBW involves modeling the conditional means of post-treatment confounders instead of the conditional distributions of the treatment to construct the weights. RBW is thus easier to use with continuous treatments, and the method is less susceptible to model misspecification issues that often arise when modeling the conditional distributions of treatments. RBW is also advantageous from a computational perspective. Because its weighting procedure involves a convex optimization problem, RBW typically locates a solution considerably faster than other methods whose optimization relies on nonconvex loss functions — such as the recently proposed nonparametric version of CBPS. We explain the rationale behind RBW, describe the functions in **rbw**, and then use real-world data to illustrate their applications in three scenarios: effect estimation for point treatments, causal mediation analysis, and effect estimation for time-varying treatments with time-varying confounders.

## **1 Introduction**

This paper describes the R package **rbw**, which implements the method of residual balancing weights for estimating marginal structural models (MSMs) (Zhou and Wodtke 2020). MSMs seek to estimate causal effects in the presence of post-treatment confounding — a common issue in the social sciences. In studies of time-varying treatments, prior treatments may affect the confounders of future treatments. For example, research has shown that political candidates’ decision to run negative advertisements is determined by their position in recent polling data, which are in turn affected by their previous decisions to run negative advertisements (Lau, Sigelman, and Rovner 2007; Blackwell 2013). Post-treatment confounding can also occur in causal mediation analysis when confounders of the mediator-outcome relationship are affected by the treatment. For example, such a problem arises in a study of the mediating role of morality in the effect of shared democracy on public support for war (Tomz and Weeks 2013). Post-treatment variables, such as respondents’ beliefs about the likelihood of victory, may determine both perceptions of morality and support for military intervention.

MSMs aim to address two types of bias associated with conventional regression methods that adjust naively for post-treatment confounders: overcontrol and collider-stratification bias (Robins 1986, 2000). Conditioning naively on post-treatment confounders can create overcontrol bias as it blocks the effect of the treatment on the outcome that passes through these variables. Additionally, conditioning naively on post-treatment confounders can lead to collider-stratification bias when these variables are affected by unobserved determinants of the outcome. This is because the adjustment will create a spurious association between the treatment and the unobserved variables.

Researchers often use inverse probability weighting (IPW) to fit MSMs (for an in-depth explanation of the method, see Robins, Hernan, and Brumback 2000; Robins 2000; Cole and Hernán 2008). In longitudinal settings, MSMs with IPW involve fitting a model for the conditional mean of the outcome given the treatment history using weights that break the dependence between past confounders and the treatment at each time point. In essence, the weights create a pseudo-population where the marginal mean of the potential outcomes

under a treatment history equals the conditional mean of the observed outcome given the treatment history. The R package `ipw` provides functions for estimating inverse probability weights (van der Wal and Geskus 2011; Geskus and van der Wal 2015).

However, IPW’s success depends on correct specification of the models for the conditional distributions of exposure to treatment and/or mediator (hereafter jointly referred to as “exposures”), which is difficult to achieve in practice. Moreover, even when these models are correctly specified, IPW is inefficient and susceptible to finite-sample bias (Zhou and Wodtke 2020; Wang et al. 2006). Finally, when the exposures are continuous, IPW may perform poorly due to unreliable estimation of conditional densities (Naimi et al. 2014; Vansteelandt 2009).

Alternative methods have attempted to mitigate these shortcomings. In particular, Imai and Ratkovic’s (2014, 2015) covariate balancing propensity score (CBPS) method proposes a set of balancing conditions when estimating the propensity scores. Because it seeks to maximize the covariate balance between the treatment and control groups, this method is less sensitive to model misspecification than IPW. Fong, Hazlett, and Imai (2018) expand CBPS to accommodate continuous exposures, but the challenges involved in modeling conditional densities persist. With this in mind, they have also developed a nonparametric extension that constructs weights that maximize the empirical likelihood while meeting a set of balancing conditions. Though the nonparametric CBPS (npCBPS) circumvents the need for specifying a functional form for the propensity score, it does so at a cost — since the empirical likelihood is not generally convex, the optimization procedure is often slow and may fail to find a solution. The latter can happen, for example, when we have a large number of covariates. The authors advance a workaround that adds flexibility to the covariate balancing conditions and penalizes the remaining imbalance. In doing so, they ensure that a weighting solution exists. Users can implement CBPS in R with the `CBPS` package (Fong, Ratkovic, and Imai 2021).

Recently, Zhou and Wodtke (2020) propose the method of residual balancing weights (RBW) for estimating MSMs. RBW involves fitting models for the conditional means of post-treatment confounders given past treatments and confounders and extracting their residual terms. It then uses Hainmueller’s (2012) entropy balancing method to find weights such that, in the weighted sample, 1) the residuals are orthogonal to future exposures, past treatments, and past confounders, and 2) the relative entropy between the weights and a set of base weights (e.g., survey sampling weights) is minimized. RBW is similar to npCBPS in that it relies on a set of balancing conditions to find the weights and does not require modeling the conditional distributions of the exposures. Both methods can, therefore, be easily adapted to cases where exposures are continuous. Despite their similarities, RBW has a significant computational advantage: the relative entropy metric it uses to construct the weights leads to a convex optimization problem, so finding the weighting solution is computationally expeditious. As shown below, RBW manages to locate the solution considerably faster than npCBPS when we compare the methods’ performance for the same problem.

In the sections that follow, we present an overview of the residual balancing method (Section 2) and how, in addition to contexts involving time-varying treatments, we can use it in cases of point treatments and causal mediation analysis (Section 3). We then discuss the package that implements the method in R (`rbw`). Section 4 describes the functions included in `rbw`, and Section 5 illustrates their use with various data sets. Section 6 concludes.

## 2 Overview of Residual Balancing

This section gives an overview of RBW (see Zhou and Wodtke 2020 for a thorough exposition of the method). We first describe the notation used throughout the paper and briefly review MSMs. Next, we explain the underlying logic of RBW and provide an intuition for how the method works using a directed acyclic graph (DAG).

## 2.1 Notation

Assume we have a study with  $T \geq 2$  time points, and we are interested in the effect of a time-varying treatment,  $A_t$  ( $1 \leq t \leq T$ ), on some end-of-study outcome  $Y$ . We also have a vector of observed time-varying confounders,  $L_t$ , at each time point, which may be affected by prior treatments.  $\bar{A}_t = (A_1, \dots, A_t)$  and  $\bar{L}_t = (L_1, \dots, L_t)$  denote treatment and covariate histories up to time  $t$ . Furthermore,  $\bar{A} = \bar{A}_T$  and  $\bar{L} = \bar{L}_T$  represent a respondent's complete treatment and covariate histories, respectively. Finally, let  $Y(\bar{a})$  be the potential outcome under some treatment history  $\bar{a}$ .

## 2.2 MSMs

An MSM is a model for the marginal mean of the potential outcomes under some treatment history:

$$\mathbb{E}[Y(\bar{a})] = \mu(\bar{a}; \beta), \quad (1)$$

where  $\mu(\cdot)$  is some function and  $\beta$  are a set of parameters capturing the causal effects of interest. We can identify an MSM from observed data under three assumptions:

1. consistency: if  $\bar{A} = \bar{a}$ , then  $Y = Y(\bar{a})$ ;
2. sequential ignorability: at each time point  $t$ , treatment is unconfounded conditional on past treatments and the covariate history up to that point. Formally,  $Y(\bar{a}) \perp\!\!\!\perp A_t | \bar{A}_{t-1}, \bar{L}_t$ ; and
3. positivity: at each time point  $t$ , treatment assignment must not be deterministic. That is, if  $f(\bar{A}_{t-1} = \bar{a}_{t-1}, \bar{L}_t = \bar{l}_t) > 0$ , then  $f(A_t = a_t | \bar{A}_{t-1} = \bar{a}_{t-1}, \bar{L}_t = \bar{l}_t) > 0$ , where  $f(\cdot)$  represents a probability mass or density function.

Under these assumptions, Robins (1986) shows that the expected value of the potential outcome  $\mathbb{E}[Y(\bar{a})]$  can be identified via the g-computation formula:

$$\mathbb{E}[Y(\bar{a})] = \int \dots \int \mathbb{E}[Y | \bar{A} = \bar{a}, \bar{L} = \bar{l}] \prod_{t=1}^T f(l_t | \bar{l}_{t-1}, \bar{a}_{t-1}) d\mu(l_t), \quad (2)$$

where  $\mu(\cdot)$  is an appropriate dominating measure. While Equation 2 provides a general formula for identifying causal effects in the presence of time-varying treatments, directly evaluating it is often impractical, particularly when we have many covariates and time periods.

## 2.3 The Rationale Behind Residual Balancing

Now consider the formula for the conditional mean of the observed outcome  $Y$  given some treatment history:

$$\mathbb{E}[Y | \bar{A} = \bar{a}] = \int \dots \int \mathbb{E}[Y | \bar{A} = \bar{a}, \bar{L} = \bar{l}] \prod_{t=1}^T f(l_t | \bar{l}_{t-1}, \bar{a}) d\mu(l_t). \quad (3)$$

By comparing Equations 2 and 3, we see that weighting the observed population by

$$W_l = \prod_{t=1}^T \frac{f(L_t | \bar{L}_{t-1}, \bar{A}_{t-1})}{f(L_t | \bar{L}_{t-1}, \bar{A})} \quad (4)$$

creates a pseudo-population in which  $f^*(l_t | \bar{l}_{t-1}, \bar{a}) = f^*(l_t | \bar{l}_{t-1}, \bar{a}_{t-1}) = f(l_t | \bar{l}_{t-1}, \bar{a}_{t-1})$  and  $\mathbb{E}^*[Y | \bar{A} = \bar{a}] = \mathbb{E}^*[Y(\bar{a})] = \mathbb{E}[Y(\bar{a})]$ , where  $*$  represents quantities in the pseudo-population. Estimating the conditional densities of Equation 4 is challenging because  $L_t$  is often high-dimensional.

Zhou and Wodtke (2020) demonstrate that the condition  $f^*(l_t|\bar{l}_{t-1}, \bar{a}) = f^*(l_t|\bar{l}_{t-1}, \bar{a}_{t-1}) = f(l_t|\bar{l}_{t-1}, \bar{a}_{t-1})$  implies a series of moment conditions in the pseudo-population. Most importantly,

$$\mathbb{E}^*[\delta(g(L_t))h(\bar{L}_{t-1}, \bar{A})] = \mathbb{E}^*[\delta(g(L_t))]\mathbb{E}^*[h(\bar{L}_{t-1}, \bar{A})] = 0, \quad (5)$$

where:

- $g(\cdot)$  and  $h(\cdot)$  are scalar functions.
- $\delta(g(L_t))$  is the residual of  $g(L_t)$  with respect to its conditional mean given the observed past:  $\delta(g(L_t)) = g(L_t) - \mathbb{E}[g(L_t)|\bar{L}_{t-1}, \bar{A}_{t-1}]$ .

Residual balancing aims to emulate the moment conditions (5) that would hold in the pseudo-population if it were possible to weight the observed population by  $W_l$ . To do so, the method 1) specifies a set of  $g(\cdot)$  functions,  $G(L_t) = \{g_1(L_t), \dots, g_{J_t}(L_t)\}$  and a set of  $h(\cdot)$  functions,  $H(\bar{L}_{t-1}, \bar{A}) = \{h_1(\bar{L}_{t-1}, \bar{A}), \dots, h_{K_t}(\bar{L}_{t-1}, \bar{A})\}$ ; 2) computes a set of residual terms  $\delta(g(L_t)) = g(L_t) - \mathbb{E}[g(L_t)|\bar{L}_{t-1}, \bar{A}_{t-1}]$ ; and 3) finds a set of weights such that, for any  $j, k$ , and  $t$ , the cross-moment of  $\delta(g_j(l_{jt}))$  and  $h_k(\bar{l}_{i,t-1}, \bar{a}_i)$  is zero in the weighted data. That is, RBW locates the  $rbw_i$  weights subject to the following balancing conditions:

$$\sum_{i=1}^n rbw_i c_{ir} = 0, \quad 1 \leq r \leq n_c, \quad (6)$$

where  $c_{ir}$  is the  $r$ th element of  $\mathbf{c}_i = \{\delta(g_j(l_{jt}))h_k(\bar{l}_{i,t-1}, \bar{a}_i); 1 \leq j \leq J_t, 1 \leq k \leq K_t, 1 \leq t \leq T\}$  and  $n_c = \sum_{t=1}^T J_t K_t$  denotes the total number of balancing conditions. The residualized confounders at each time point are balanced across future treatments as well as past treatments and confounders (the observed past). RBW thus adjusts for post-treatment confounding without inducing overcontrol and collider-stratification bias.

Moreover, Zhou and Wodtke (2020) follow Hainmueller (2012) and minimize the relative entropy between  $rbw_i$  and a set of base weights  $q_i$  (e.g., survey sampling weights):

$$\min_{rbw_i} \sum_i rbw_i \log(rbw_i/q_i) \quad (7)$$

We can then use the method of Lagrange multipliers to find a weighting solution that minimizes the relative entropy between  $rbw_i$  and  $q_i$  subject to the  $n_c$  balancing constraints. We discuss this procedure in greater depth in Section 4.1 when describing the function `rbw:::eb2()`. The convexity of the relative entropy metric renders it considerably more computationally efficient than nonconvex loss functions that can also be used to construct weights, such as the empirical likelihood (Fong, Hazlett, and Imai 2018).

A typical implementation of residual balancing can be summarized in three steps:

1. For each covariate  $j$  and at each time point  $t$ , fit a linear, logistic, or Poisson regression of  $l_{ijt}$  (depending on its level of measurement) on  $\bar{l}_{i,t-1}$  and  $\bar{a}_{i,t-1}$ . Then compute the residuals  $\hat{\delta}(l_{ijt})$ . For the covariates in  $L_1$  (the first time period), the residuals are computed as deviations from the sample mean:  $\hat{\delta}(l_{ij1}) = l_{ij1} - \text{avg}(l_{j1})$ . This step relies on the idea that  $g_j(L_t) = L_{jt}$ , where  $L_{jt}$  is the  $j$ th element of the covariate vector  $L_t$ , is a natural choice for the set of  $g(\cdot)$  functions constituting  $G(L_t)$ .
2. Find a set of weights,  $rbw_i$ , such that:
  - a) in the weighted sample, the residuals  $\hat{\delta}(l_{ijt})$  are orthogonal to all future treatments and the regressors of  $l_{ijt}$  (i.e., the past treatments and past confounders);
  - b) the relative entropy between  $rbw_i$  and the base weights  $q_i$  is minimized.
3. Use the weights to fit an MSM.

Figure 1 depicts the logic of RBW in a DAG. Following the notation described above,  $A_t$  denotes our time-varying treatment,  $L_t$  is a vector of time-varying confounders, and  $Y$  is our end-of-study outcome. Further, assume two time points,  $t = 1, 2$ . The  $rbw_i$  weights break the dependence between  $A_t$  and  $\bar{L}_t$  at each time point. That is, the weights create a pseudo-population where the confounding arrows  $L_1 \rightarrow A_1$ ,  $L_1 \rightarrow A_2$ , and  $L_2 \rightarrow A_2$  are broken while all others are preserved. It is important to note that  $L_1$  is *marginally* independent of both  $A_1$  and  $A_2$  in the pseudo-population, but  $L_2$  is *conditionally* independent of  $A_2$ , given  $L_1$  and  $A_1$ . Hence, RBW invokes a model for the conditional mean of  $L_2$  given  $L_1$  and  $A_1$  and balances the residuals from this model across levels of  $A_2$  and levels of  $(L_1, A_1)$  (the observed past). This procedure avoids overcontrol and collider-stratification bias when adjusting for post-treatment confounding because it breaks the confounding arrow  $L_2 \rightarrow A_2$  while leaving the causal arrow  $A_1 \rightarrow L_2$  intact.

Finally, since  $\mathbb{E}^*[Y|\bar{A} = \bar{a}] = \mathbb{E}^*[Y(\bar{a})] = \mathbb{E}[Y(\bar{a})]$  in the pseudo-population, we can estimate the marginal effects of interest by fitting a model for the conditional mean of the observed outcome given the treatment history (and possibly a set of baseline confounders) with weights equal to  $rbw_i$ .




Figure 1: The underlying logic of residual balancing.

### 3 Uses of Residual Balancing

The rationale described in the previous section is targeted at estimating the causal effects of time-varying treatments. With minor adaptations, we can expand the use of RBW to two other contexts commonly encountered in the social and biomedical sciences: point treatment situations and causal mediation analysis.

#### 3.1 RBW for Estimating the Average Effect of a Point Treatment

RBW can be easily adapted to a point treatment situation where the user aims only to adjust for a set of baseline (i.e., time-invariant) confounders to estimate the average treatment effect. To this end, we modify the procedure above as follows:

1. Compute the response residuals  $\hat{\delta}(x_{ij})$  for each baseline confounder  $X_j$  by centering it around its sample mean:  $\hat{\delta}(x_{ij}) = x_{ij} - \text{avg}(x_j)$ .
2. Find a set of weights,  $rbw_i$ , such that:
  - a) in the weighted sample, the residuals  $\hat{\delta}(x_{ij})$  are orthogonal to the treatment;
  - b) the relative entropy between  $rbw_i$  and the base weights  $q_i$  is minimized.
3. Use the weights to fit an MSM.

The DAG of Figure 2 illustrates the point treatment situation.  $A$  represents the one-shot treatment,  $X$  is a vector of baseline confounders, and  $Y$  denotes our outcome. Weighting the observed population by  $rbw_i$  mimics a pseudo-population where the link between  $A$  and  $X$  is broken. We can then fit a model for the conditional mean of  $Y$  given  $A$  to estimate the causal effects of interest.<sup>1</sup>




Figure 2: RBW in a Point Treatment.

### 3.2 RBW in Causal Mediation Analysis

In causal mediation analysis, researchers are often concerned with the joint effects of a one-shot treatment,  $A$ , and a mediator,  $M$ , on some end-of-study outcome  $Y$  when both baseline confounders ( $X$ ) and some post-treatment confounders for the mediator-outcome relationship ( $Z$ ) are present. With minor adjustments, the RBW implementation for causal mediation analysis is similar to the case with time-varying treatments:

1. As in the point treatment scenario, compute the response residuals  $\hat{\delta}(x_{ij})$  for each baseline confounder  $X_j$  by centering it around its sample mean.
  - Note: users may skip this step while including the baseline confounders in the MSM in the final step.
2. Estimate the response residuals  $\hat{\delta}(z_{ij})$  for each post-treatment confounder  $Z_j$  by fitting a linear, logistic, or Poisson regression of  $z_{ij}$  (depending on its level of measurement) on the treatment  $a_i$  and the baseline confounders  $x_i$ :  $\hat{\delta}(z_{ij}) = z_{ij} - \mathbb{E}[z_{ij}|a_i, x_i]$ .
3. Find a set of weights,  $rbw_i$ , such that:
  - a) in the weighted sample, i) the baseline residuals  $\hat{\delta}(x_{ij})$  are orthogonal to the treatment  $a_i$  and the mediator  $m_i$ ; and ii) the post-treatment residuals  $\hat{\delta}(z_{ij})$  are balanced across the treatment  $a_i$ , the mediator  $m_i$ , and the baseline confounders  $x_i$ ;
  - b) the relative entropy between  $rbw_i$  and the base weights  $q_i$  is minimized.

<sup>1</sup>The vector  $X$  captures all confounding under the ignorability assumption.

4. Use the weights to fit an MSM for the joint effects of the treatment and the mediator on the outcome:
  - a) In causal mediation analysis, the potential outcomes of interest are denoted by  $Y(a, m)$  (this is the potential outcome under treatment  $a$  and mediator value  $m$ ). We can then express a saturated MSM as  $\mathbb{E}[Y(a, m)] = \alpha_0 + \alpha_1 a + \alpha_2 m + \alpha_3 am$ .
  - b) Alternatively, the baseline confounders can be included in the MSM if users decide to skip the first step:  $\mathbb{E}[Y(a, m)|X] = \alpha_0 + \alpha_1 a + \alpha_2 m + \alpha_3 am + \alpha_4^T X$ .
  - c) Finally, the controlled direct effects (CDE) of the treatment can be estimated as  $\widehat{\text{CDE}}(m) = \mathbb{E}[Y(1, m) - Y(0, m)] = \hat{\alpha}_1 + \hat{\alpha}_3 m$ . The CDE measures the causal effect of the treatment on the outcome when the mediator is fixed at some value  $m$  for all units.

## 4 The R package

The R package `ipw` contains four functions:

- `eb2()`, for generating minimum entropy weights subject to a set of balancing constraints.
- `rbwATE()`, for constructing residual balancing weights to estimate the causal effects of one-shot treatments.
- `rbwMed()`, for constructing residual balancing weights to estimate controlled direct effects in causal mediation analysis.
- `rbwPanel()`, for constructing residual balancing weights to estimate the marginal effects of time-varying treatments.

Next, we explain each of these functions. The package also includes several real-world data sets (`advertisement`, `peace`, `campaign_long`, and `campaign_wide`), which we describe and analyze in the examples of Section 5.

### 4.1 Function `eb2()`

This function is an adaptation of `ebal::eb()` (Hainmueller 2014). It is called internally by other functions in `rbw` to implement the method of Lagrange multipliers for locating a weighting solution that minimizes the relative entropy between  $rbw_i$  and  $q_i$  subject to the set of  $n_c$  balancing constraints described in Equation 6. Zhou and Wodtke (2020) impose an additional normalization constraint that ensures that the residual balancing weights sum to the sample size:  $\sum_i rbw_i = n$ . Following Hainmueller (2012), the authors obtain the primal optimization problem:

$$\min_{rbw_i} L^p = \sum_{i=1}^n rbw_i \log(rbw_i/q_i) + \sum_{r=1}^{n_c} \lambda_r \sum_{i=1}^n rbw_i c_{ir} + \lambda_0 \left( \sum_{i=1}^n rbw_i - n \right), \quad (8)$$

where  $\{\lambda_1, \dots, \lambda_{n_c}\}$  are the Lagrange multipliers for the balancing constraints and  $\lambda_0$  is the Lagrange multiplier for the normalization constraint. Since the loss function  $L^p$  is strictly convex, the first order condition of  $\frac{\partial L^p}{\partial rbw_i} = 0$  implies that the solution for each weight is

$$rbw_i^* = \frac{nq_i \exp(-\sum_{r=1}^{n_c} \lambda_r c_{ir})}{\sum_{i=1}^n q_i \exp(-\sum_{r=1}^{n_c} \lambda_r c_{ir})}. \quad (9)$$

We can then insert Equation 9 into  $L^p$ , leading to an unrestricted dual problem:

$$\max_{\lambda_r} L^d = -\log \left( \sum_{i=1}^n q_i \exp(-\sum_{r=1}^{n_c} \lambda_r c_{ir}) \right), \quad (10)$$

or equivalently,

$$\min_Z L^d = \log \left( Q' \exp(CZ) \right), \quad (11)$$

where  $Q = [q_1, \dots, q_n]'$ ,  $C = [c_1, \dots, c_1]'$ , and  $Z = -[\lambda_1, \dots, \lambda_{n_c}]'$ . Since  $L^d$  is strictly convex, the solution is guaranteed to be unique — assuming one exists. Given that both the gradient and the Hessian have closed-form expressions, we can solve the problem using Newton’s method. `eb2()` implements the algorithm. If convergence is successful, the function tells the user that “Entropy minimization converged within tolerance level.” Otherwise, it warns that entropy minimization did not converge and suggests increasing the number of iterations or reducing the number of balancing constraints.

The convexity of our optimization problem leads to appreciable computational gains over other methods that use alternative loss functions. This will be demonstrated in Section 5.1, when we compare the performance of RBW with that of npCBPS, which uses the empirical likelihood, for the same problem.

`eb2()` is used as:

```
eb2(C, M, Q, Z = rep(0, ncol(C)), max_iter = 200, tol = 1e-04,
    print_level = 2)
```

and takes the following arguments:

- `C` is a constraint matrix, with each column representing a balancing constraint.
- `M` is a vector of moment conditions to be met in the reweighted sample (per Equation 6, this is a vector of zeros with length equal to the number of columns of `C` when the other functions in `rbw` call `eb2()` internally).
- `Q` is a vector of base weights.
- `Z` is a vector of Lagrange multipliers to be initialized.
- `max_iter` determines the maximum number of iterations for the Newton’s method.
- `tol` is a tolerance parameter used to determine the convergence of the Newton’s method.
- `print_level` determines the level of printing:
  - **1** normal: print whether the algorithm converges or not
  - **2** detailed: print also the maximum absolute value of the deviation between the moments of the reweighted data and the target moments in each iteration
  - **3** very detailed: print also the step length of the line searcher in iterations where a full Newton step is excessive

The output returned by `eb2()` is a list containing the following elements:

- `W` is a vector of normalized minimum entropy weights.
- `Z` is a vector of Lagrange multipliers.
- `converged` is a logical indicator for convergence.
- `maxdiff` is a scalar indicating the maximum absolute value of the deviation between the moments of the reweighted data and the target moments in each iteration.

## 4.2 Function `rbwATE()`

This function produces residual balancing weights to be used in a point treatment situation. It implements the steps described in Section 3.1. It first takes a set of baseline confounders and computes the response residuals for each confounder by centering it around its sample mean. Then it calls `eb2()` to find a set of weights,  $rbw_i$ , such that 1) the baseline residuals are orthogonal to the treatment in the weighted sample, and 2) the relative entropy between  $rbw_i$  and the base weights is minimized. Additionally, `rbwATE()` calls a function that ensures that the matrix of balancing constraints comprises only linearly independent columns.



`rbwATE()` is used as:

```
rbwATE(treatment, data, baseline_x, base_weights, max_iter = 200,  
       print_level = 1, tol = 1e-6)
```

and takes the following arguments:

- `max_iter`, `print_level`, and `tol` have the same definitions as in `eb2()`.
- `treatment` is a symbol or character string for the treatment variable.
- `data` is a data frame containing all variables in the model.
- `baseline_x` is an expression for a set of baseline confounders stored in `data`.
- `base_weights` is an *optional* vector of base weights (or its name). If no value is supplied, the function sets a vectors of ones with length equal to the sample size as the base weights.

The output returned by `rbwATE()` is a list containing the following elements:

- `weights` is a vector of residual balancing weights.
- `constraints` is a matrix of linearly independent residual balancing constraints.
- `eb_out` contains the results from calling the `eb2()` function.
- `call` is the matched call (the function call with all arguments specified by their full names).

### 4.3 Function `rbwMed()`

This function produces residual balancing weights for causal mediation analysis. It implements the steps described in Section 3.2 by taking an *optional* set of baseline confounders (as explained above, users can opt instead to include these covariates in the MSM later) and a list of model objects for the conditional mean of the each post-treatment confounder given the treatment and the baseline confounders. It then calls `eb2()` to find a set of weights,  $rbw_i$ , such that, in the weighted sample, 1) the baseline residuals are orthogonal to the treatment and the mediator; 2) the post-treatment confounders are balanced across the treatment, the mediator, and the baseline confounders; and 3) the relative entropy between  $rbw_i$  and the base weights is minimized.

`rbwMed()` takes an additional argument, `interact`, a logical variable indicating whether the baseline and post-treatment confounders should be balanced against the treatment-mediator interaction. This argument is set to `FALSE` by default, but users suspecting an interaction effect may find it prudent to balance against it. Like `rbwATE()`, `rbwMed()` calls a function internally to ensure that only linearly independent columns constitute the matrix of balancing constraints.

It is used as:

```
rbwMed(treatment, mediator, zmodels, data, baseline_x, interact = FALSE,  
       base_weights, max_iter = 200, print_level = 1, tol = 1e-04)
```

and takes the following arguments:

- `treatment`, `data`, `base_weights`, `max_iter`, `print_level`, and `tol` are defined as in `rbwATE()`.
- `baseline_x` is an *optional* expression for a set of baseline confounders stored in `data`.
- `mediator` is a symbol or character string representing the mediator variable.
- `zmodels` is a list of fitted `lm` or `glm` objects for post-treatment confounders of the mediator-outcome relationship. Users should set this argument to `NULL` if there are no post-treatment confounders.
- `interact` is a logical variable indicating whether baseline and post-treatment confounders should be balanced against the treatment-mediator interaction term(s).

The output returned by `rbwMed()` is a list containing the same elements as the output from `rbwATE()`.

## 4.4 Function `rbwPanel()`

This function produces residual balancing weights for estimating the marginal effects of time-varying treatments. It implements the steps described in Section 2.3 by taking a list of model objects for the conditional mean of each post-treatment confounder given past treatments and past confounders. Then it calls `eb2()` to find a set of weights,  $rbw_i$ , such that 1) residuals of the post-treatment confounders are orthogonal to future treatments and the observed past in the weighted sample, and 2) the relative entropy between  $rbw_i$  and the base weights is minimized. Like the other functions, `rbwPanel()` ensures that the matrix of balancing constraints consists only of linearly independent columns.

It is used as:

```
rbwPanel(treatment, xmodels, id, time, data, base_weights, max_iter = 200,
         print_level = 1, tol = 1e-04, future = 1L)
```

and takes the following arguments:

- `data`, `base_weights`, `max_iter`, `print_level`, and `tol` are defined as in `rbwATE()` and `rbwMed()`.
- `treatment` is a symbol or character string for the time-varying treatment.
- `xmodels` is a list of fitted `lm` or `glm` objects for time-varying confounders.
- `id` is a symbol or character string for the unit id variable.
- `time` is a symbol or character string for the time variable.
- `future` is an integer indicating the number of future treatments in the balancing conditions. When `future > 0`, the residualized time-varying covariates are balanced not only with respect to current treatment  $A_t$ , but also with respect to future treatments  $A_{t+1}, \dots, A_{t+\text{future}}$ . The default, `future = 1`, assumes away higher-ordered lagged effects of the covariates on the treatment. Users can leave out lagged effects entirely by setting `future` to zero.

The output returned by `rbwPanel()` is essentially the same as those from `rbwATE()` and `rbwMed()`. The only difference is that the `weights` object, instead of a vector, is now a data frame with two columns: the id variable and the residual balancing weights (the column storing the weights is called `rbw`).

## 5 Examples

We now illustrate the functions `rbwPanel()`, `rbwATE()`, and `rbwMed()` with data sets `advertisement`, `peace`, `campaign_long`, and `campaign_wide`, which are included in `rbw`. We expect users to rarely need to call `eb2()` manually since its primary use is to be called internally by the other functions. Hence, we see little gain in providing a separate example for it.

### 5.1 Point Treatment: Effects of Political Advertisements on Campaign Contributions

Urban and Niebler (2014) studied the effects of televised political advertisements on campaign contributions. Presidential candidates do not tend to deliberately advertise in states where competition for electoral votes is tame. Yet, some areas of noncompetitive states have overlapping media markets with battleground states. Because these media market spillovers do not encompass other forms of campaigning (e.g., rallies, speeches, etc.), the authors can isolate the effect of television advertising by restricting their analyses to noncompetitive states. Their original method involved estimating the propensity score with logistic regression and then conducting propensity score matching. To do so, they dichotomized the political advertising variable to indicate whether zip codes received more than 1000 advertisements.

Deeming this approach inadequate — in part because balancing against a dichotomous treatment does not ensure covariate balance on the underlying continuous variable — Fong, Hazlett, and Imai (2018) replicate

the study using the CBPS method applied to a continuous treatment. Next, we examine how RBW fares compared with CBPS in this point treatment situation.<sup>2</sup>

Importantly, CBPS assumes that the treatment variable is normally distributed. To satisfy this assumption, Fong, Hazlett, and Imai (2018) search across Box-Cox transformations of the form

$$a_i^* = \begin{cases} \frac{(a_i+1)^\lambda - 1}{\lambda} & \text{if } \lambda \neq 0, \\ \log(a_i + 1) & \text{if } \lambda = 0 \end{cases} \quad (12)$$

to find the most appropriate transformation of the treatment.  $a_i$  is the original treatment variable, the total number of political advertisements in a zip code, and  $a_i^*$  is the transformed treatment. The authors add one to  $a_i$  inside the Box-Cox transformations to avoid  $\log(0)$  for zip codes that received no advertisements. They find the transformation with  $\lambda = -0.16$  under which the marginal distribution of the treatment is closest to the standard normal.

Let  $A^*$  denote the Box-Cox transformed treatment. We also have a vector  $X$  of baseline confounders that include the log population, population density, log median income, percent Hispanic, percent black, percent over age 65, percent college graduates, and a binary indicator of whether it is possible to commute to the zip code from a competitive state. Finally,  $Y$  represents the outcome, campaign contributions in thousands of dollars.

Fong, Hazlett, and Imai (2018) propose an MSM that includes the transformed treatment variable, its square, and state dummies  $U$  to account for state fixed effects:

$$\mathbb{E}[Y(a^*)|U] = \theta_0 + \theta_1 a^* + \theta_2 (a^*)^2 + \theta_3^T U. \quad (13)$$

The data set `advertisement` contains the variables necessary for the analyses:

- `zip` denotes zip code.
- `TotAds` is the total number of political advertisements aired in the zip code.
- `treat` is the Box-Cox transformed `TotAds`.
- `TotalPop` is the population size.
- `log_TotalPop` is the log population size
- `Inc` is the median household income.
- `log_Inc` is the log median household income.
- `PercentOver65` is the percent of the population over 65.
- `PercentHispanic` is the percent of the population Hispanic.
- `PercentBlack` is the percent of the population black.
- `per_collegegrads` is the percent college graduates.
- `density` is the population density (people per square mile).
- `CanCommute` is a dummy variable indicating whether it is possible to commute to the zip code from a competitive state.
- `Cont` measures campaign contributions (in thousands of dollars).
- `StFIPS` is the FIPS code of the state.

We start by loading the necessary packages:

```
library("rbw")
library("CBPS")
```

<sup>2</sup>Readers can find examples for IPW and its implementation using the R package `ipw` in van der Wal and Geskus (2011). Fong, Hazlett, and Imai (2018) and Zhou and Wodtke (2020) use extensive simulation studies to demonstrate how their methods are more efficient and more accurate than IPW, so we do not draw similar comparisons here.

`rbw::rbwATE()` will construct the residual balancing weights by following the steps described in Section 3.1. It first computes the baseline residuals  $\hat{\delta}(x_{ij}) = x_{ij} - \text{avg}(x_j)$  and then finds a set of weights such that, in the weighted sample,  $\hat{\delta}(x_{ij})$  are orthogonal to the treatment, and the relative entropy between the residual balancing weights and a set of base weights is minimized. Since `advertisement` does not include sampling weights, `rbw::rbwATE()` will set a vector of ones with length equal to the sample size as the base weights. To facilitate comparisons of computation time, we also call `Sys.time()` before and after running the functions that produce the weights in each package.

```
rbwATE_start <- Sys.time()
rbwATE_fit <- rbwATE(treat, baseline_x = c(log_TotalPop, PercentOver65,
  log_Inc, PercentHispanic, PercentBlack, density,
  per_collegegrads, CanCommute), data = advertisement)
```

```
## Entropy minimization converged within tolerance level
```

```
rbwATE_end <- Sys.time()
rbwATE_time <- rbwATE_end - rbwATE_start
```

Next, we attach the  $rbw_i$  weights to the data:

```
advertisement$rbwATE_weights <- rbwATE_fit$weights
```

Following most applications of MSMs, we compute standard errors using the robust (“sandwich”) variance estimator. This can be implemented with the function `survey::svydesign()`, which allows us to specify a complex survey design and estimate standard errors consistent with this specification.<sup>3</sup>

```
library("survey")
```

```
## Warning: package 'survey' was built under R version 3.6.3
```

```
rbwATE_design <- svydesign(ids = ~1, weights = ~rbwATE_weights,
  data = advertisement)
```

We then use the residual balancing weights to fit the MSM defined in Equation 13:

```
rbwATE_msm <- svyglm(Cont ~ treat + I(treat^2) + factor(StFIPS),
  design = rbwATE_design)
```

Since we have transformed our treatment variable, we need to make necessary adjustments to compute the treatment effect. Let us first use a scalar object to store the parameter estimated from the Box-Cox transformation:

```
lambda <- -0.16
```

Like Urban and Niebler (2014), Fong, Hazlett, and Imai (2018) are interested in the effect of going from zero to 1,000 political advertisements on campaign contributions. Hence, we create a `dose` variable to account for this. If  $A^*$  is the transformed treatment and  $A$  is the treatment variable on its original scale,  $A^* = [(A + 1)^\lambda - 1]/\lambda$ . Our dose is thus given by  $dose = [(1000 + 1)^\lambda - 1]/\lambda$ :

---

<sup>3</sup>Zhou and Wodtke (2020) conduct various simulation studies to assess the performance of the robust variance estimator across different methods. They find that the estimator tends to overestimate the true sampling variance for residual balancing across nearly all scenarios, making it consistently conservative for RBW. By contrast, the estimator sometimes overestimates and other times underestimates the true sampling variance for other weighting methods, including CBPS.

```
dose <- ((1000 + 1)^lambda - 1)/lambda
```

To find the estimate  $\hat{\tau}_{rbw}$  of the average treatment effect, we must consider both the linear and squared treatment terms in the MSM. We also multiply the result by 1,000 since the outcome variable is measured in thousands of dollars:

```
rbwATE_tau <- 1000 * (coef(rbwATE_msm)[2] * dose + coef(rbwATE_msm)[3] * dose^2)
```

We can use the basic properties of the variance operator to derive the formula for  $\text{Var}[\hat{\tau}_{rbw}]$ :

$$\text{Var}[\hat{\tau}_{rbw}] = dose^2 \text{Var}[\hat{\theta}_1] + dose^4 \text{Var}[\hat{\theta}_2] + 2 * dose^3 \text{Cov}[\hat{\theta}_1, \hat{\theta}_2]. \quad (14)$$

The function `stats::vcov()` applied to the object `rbwATE_msm` yields estimates of  $\text{Var}[\hat{\theta}_1]$ ,  $\text{Var}[\hat{\theta}_2]$ , and  $\text{Cov}[\hat{\theta}_1, \hat{\theta}_2]$ . Again, we multiply the result by 1,000 given the scale of the outcome variable:

```
rbwATE_vcov <- stats::vcov(rbwATE_msm)
rbwATE_se <- 1000 * sqrt(dose^2 * rbwATE_vcov[2, 2] +
  dose^4 * rbwATE_vcov[3, 3] +
  2 * dose^3 * rbwATE_vcov[2, 3])
```

We now repeat the steps above using the `CBPS()` and `npCBPS()` functions from the `CBPS` package to estimate weights using both the parametric and nonparametric CBPS approaches. Recall from Section 1 that the former is similar to IPW in that it requires explicit models for the conditional distributions of the treatment. However, it improves the IPW method by considering a set of balancing conditions during the propensity score estimation, thereby reducing sensitivity to model misspecification. By contrast, the nonparametric CBPS (`npCBPS`) does not require direct estimation of the propensity score; instead, it finds weights that maximize the empirical likelihood while meeting a set of balancing constraints. As such, like RBW, it avoids the need to specify a propensity score model.

Fong, Hazlett, and Imai (2018) propose a treatment model that includes all baseline confounders  $X$  defined above in addition to the square terms of all continuous covariates. We define the formula accordingly:

```
CBPS_pscore_form <- treat ~ log(TotalPop) + PercentOver65 +
  log(Inc + 1) + PercentHispanic + PercentBlack +
  density + per_collegegrads + CanCommute + I(log(TotalPop)^2) +
  I(PercentOver65^2) + I(log(Inc + 1)^2) + I(PercentHispanic^2) +
  I(PercentBlack^2) + I(density^2) + I(per_collegegrads^2)
```

We first construct the CBPS weights:

```
CBPS_start <- Sys.time()
CBPS_fit <- CBPS(CBPS_pscore_form, data = advertisement,
  twostep = TRUE, method = "exact")
CBPS_end <- Sys.time()
CBPS_time <- CBPS_end - CBPS_start
```

Next, we compute the `npCBPS` weights:

```
npCBPS_start <- Sys.time()
npCBPS_fit <- npCBPS(CBPS_pscore.form, data = advertisement,
  corprior = 0.1/nrow(advertisement))
npCBPS_end <- Sys.time()
npCBPS_time <- npCBPS_end - npCBPS_start
```

We attach them to the data:

```
advertisement$CBPS_weights <- CBPS_fit$weights
advertisement$npCBPS_weights <- npCBPS_fit$weights
```

Following the steps above, we estimate  $\hat{\tau}_{CBPS}$ :

```
CBPS_design <- svydesign(ids = ~1, weights = ~ CBPS_weights,
  data = advertisement)
CBPS_msm <- svyglm(Cont ~ treat + I(treat^2) + factor(StFIPS),
  design = CBPS_design)
CBPS_tau <- 1000 * (coef(CBPS_msm)[2] * dose + coef(CBPS_msm)[3] * dose^2)
```

Similarly for  $\hat{\tau}_{npCBPS}$ :

```
npCBPS_design <- svydesign(ids = ~1, weights = ~npCBPS_weights,
  data = advertisement)
npCBPS_msm <- svyglm(Cont ~ treat + I(treat^2) + factor(StFIPS),
  design = npCBPS_design)
npCBPS_tau <- 1000 * (coef(npCBPS_msm)[2] * dose + coef(npCBPS_msm)[3] * dose^2)
```

Then we compute the CBPS standard error:

```
CBPS_vcov <- stats::vcov(CBPS_msm)
CBPS_se <- 1000 * sqrt(dose^2 * CBPS_vcov[2, 2] +
  dose^4 * CBPS_vcov[3, 3] +
  2 * dose^3 * CBPS_vcov[2, 3])
```

As well as its npCBPS counterpart:

```
npCBPS_vcov <- stats::vcov(npCBPS_msm)
npCBPS_se <- 1000 * sqrt(dose^2 * npCBPS_vcov[2, 2] +
  dose^4 * npCBPS_vcov[3, 3] +
  2 * dose^3 * npCBPS_vcov[2, 3])
```

Table 1 summarizes the findings. The methods that do not require a propensity score model, RBW and npCBPS, yield similar point estimates and standard errors. By contrast, the estimate associated with CBPS is somewhat smaller. Finally, the last column shows that different loss functions for the optimization problem can lead to stark differences in computation time. While RBW’s relative entropy metric leads to a convex optimization problem that Newton’s method can solve in less than one second, npCBPS’s algorithm takes several minutes to run.

We next illustrate the use of `rbwMed()` and `rbwPanel()`. In contrast to the point treatment example, the following sections do not draw comparisons with other existing methods. Interested readers can refer to Zhou and Wodtke (2020) for such comparisons.

## 5.2 Causal Mediation Analysis: The Controlled Direct Effect of Shared Democracy on Public Support for War

A stylized fact in political science is that democracies do not engage in war with one another. To assess the role of public opinion in keeping the peace, Tomz and Weeks (2013) designed survey experiments that

Table 1: Comparison of RBW, CBPS, and npCBPS for a Point Treatment Situation

Method	Estimate	Standard Error	Computation Time (in Seconds)
RBW	6506	3721	0.34
CBPS	4935	3551	8.23
npCBPS	6518	3521	147.15

Computation time may differ depending on system setup.

System setup used to generate the results: MacBook Pro (15-inch, 2018), 2.2 GHz 6-Core Intel Core i7, 16GB RAM.

presented participants with a situation where a country was developing nuclear weapons. When describing the situation, the authors randomly and independently changed three characteristics of the country: its political regime (whether it was a democracy), alliance status (whether it had signed a military alliance with the United States), and economic ties (whether it had high levels of trade with the U.S.). The outcome of interest was the respondent’s support for military action on a five-point scale ranging from “oppose strongly” to “favor strongly.” The authors found that respondents were considerably less supportive of military action against democracies than otherwise identical autocratic regimes.

They then went on to investigate the causal mechanisms through which shared democracy reduces public enthusiasm for war. In particular, Tomz and Weeks (2013) measured individuals’ beliefs about the level of threat posed by the potential adversary, the cost of the intervention, and the likelihood of success. They also collected data on people’s moral concerns about using military force. Their methodological framework focuses on estimating the natural direct and natural indirect effects (Imai et al. 2011; Imai, Keele, and Yamamoto 2010), whose identification assumptions require that no post-treatment confounding of the mediator-outcome relationship exists. As such, the authors examined the role of each mediator separately by assuming they operate independently of one another. Yet, as discussed in Zhou and Wodtke (2020), beliefs about the threat, cost, and likelihood of success may affect an individual’s perceptions of morality while also influencing support for war directly. By treating these variables as post-treatment confounders, Zhou and Wodtke (2020) analyzed the mediating role of morality using controlled direct effects.

Let  $A$  denote the treatment, whether the country developing nuclear weapons was presented as a democracy,  $M$  the mediator, a dummy variable indicating whether the participant deemed the military action morally wrong, and  $Y$  the outcome, the respondent’s support for war on a five-point scale. We also have a set of baseline confounders ( $X$ ) including dummies for the other two randomized treatments (alliance status and economic ties) in addition to several demographic and attitudinal controls.<sup>4</sup> Finally,  $Z$  is a vector of post-treatment confounders comprising measures of beliefs about the threat, cost, and likelihood of success.

Our MSM is thus defined as:

$$\mathbb{E}[Y(a, m)] = \alpha_0 + \alpha_1 a + \alpha_2 m + \alpha_3 am. \tag{15}$$

We can alternatively include the baseline confounders in the MSM instead of balancing them across the treatment and the mediator with baseline residuals:

$$\mathbb{E}[Y(a, m)|X] = \alpha_0 + \alpha_1 a + \alpha_2 m + \alpha_3 am + \alpha_4^T X. \tag{16}$$

The `peace` data set includes the variables we will use in the analyses:

- `democ` is a dummy variable indicating whether the country was a democracy.
- `threatc` is the number of adverse events respondents considered probable if the US did not engage in war.

<sup>4</sup>For example, attitudinal controls include respondents’ attitudes toward ethnocentrism with a series of questions about their opinions on immigration, affirmative action, and gay marriage.



- `cost` is the number of negative consequences anticipated if the US engaged in war.
- `successc` indicates whether the respondent thought the operation would succeed:
  - 0: less than 50-50 chance of working even in the short run
  - 1: efficacious only in the short run
  - 2: successful both in the short and long run
- `immoral` indicates whether respondents thought it would be morally wrong to strike the country.
- `ally` indicates whether the country had signed a military alliance with the US.
- `trade` indicates whether the country had high levels of trade with the US.
- `h1` is an index measuring respondent's attitude toward militarism.
- `i1` is an index measuring respondent's attitude toward internationalism.
- `p1` is an index measuring respondent's identification with the Republican party.
- `e1` is an index measuring respondent's attitude toward ethnocentrism.
- `r1` is an index measuring respondent's attitude toward religiosity.
- `male` indicates whether the respondent is male.
- `white` indicates whether the respondent is white.
- `age` is the respondent's age.
- `ed4` is the respondent's education with categories ranging from high school or less to postgraduate degree.
- `strike` is a measure of support for war on a five-point scale.

We first consider the approach where the baseline confounders are balanced across the treatment and the mediator using the baseline residuals computed from centering each covariate around its sample mean.

As explained in Section 3.2, RBW in causal mediation analysis requires models for the conditional means of each post-treatment confounder  $Z_j$  given the treatment  $a_i$  and the baseline confounders  $x_i$ :  $\hat{\mathbb{E}}[z_{ij}|a_i, x_i]$ . Hence, we assume that `threatc`, `cost`, and `successc` are measured on a continuous scale and fit linear models for each:

```
z1 <- lm(threatc ~ ally + trade + h1 + i1 + p1 + e1 + r1 +
  male + white + age + ed4 + democ, data = peace)
z2 <- lm(cost ~ ally + trade + h1 + i1 + p1 + e1 + r1 +
  male + white + age + ed4 + democ, data = peace)
z3 <- lm(successc ~ ally + trade + h1 + i1 + p1 + e1 + r1 +
  male + white + age + ed4 + democ, data = peace)
```

We then store the three model objects together in a list to be passed later to the `zmodels` argument in `rbwMed()`:

```
zmodels <- list(z1, z2, z3)
```

To construct the residual balancing weights, `rbwMed()` will 1) compute the baseline residuals  $\hat{\delta}(x_{ij}) = x_{ij} - \text{avg}(x_j)$  and the post-treatment residuals  $\hat{\delta}(z_{ij}) = z_{ij} - \mathbb{E}[z_{ij}|a_i, x_i]$  and 2) find a set of weights such that a) in the weighted sample, the baseline and post-treatment residuals meet the orthogonality requirements explained in Section 3.2, and b) the relative entropy between the residual balancing weights and a set of base weights is minimized. The function will use a vector of ones as the base weights since `peace` does not include sampling weights. We also pass the name of our mediator to the `mediator` argument:

```
rbwMed_fit <- rbwMed(treatment = democ, mediator = immoral,
  zmodels = zmodels, baseline_x = c(ally, trade, h1, i1,
  p1, e1, r1, male, white, age, ed4), data = peace)
```

```
## Entropy minimization converged within tolerance level
```



Next, we attach the weights to `peace`:

```
peace$rbwMed_weights <- rbwMed_fit$weights
```

We use `survey::svydesign()` to estimate robust standard errors:

```
rbwMed_design <- svydesign(ids = ~1, weights = ~rbwMed_weights,  
  data = peace)
```

Finally, we use the residual balancing weights to fit the MSM from Equation 15:

```
rbwMed_msm <- svyglm(strike ~ democ * immoral,  
  design = rbwMed_design)
```

Steps are similar for the case where the baseline confounders are not balanced against the treatment and the mediator but instead adjusted in the MSM, as defined in Equation 16. The models for the conditional means of the post-treatment confounders are the same as before:

```
rbwMed2_fit <- rbwMed(treatment = democ, mediator = immoral,  
  zmodels = zmodels, data = peace)
```

```
## Entropy minimization converged within tolerance level
```

```
peace$rbwMed2_weights <- rbwMed2_fit$weights  
rbwMed2_design <- svydesign(ids = ~ 1, weights = ~ rbwMed2_weights,  
  data = peace)  
rbwMed2_msm <- svyglm(strike ~ ally + trade + h1 + i1 + p1 +  
  e1 + r1 + male + white + age + ed4 + democ * immoral,  
  design = rbwMed2_design)
```

We summarize the results in Table 2. The estimated CDE if respondents lacked any moral qualms about military intervention, i.e.,  $\widehat{\text{CDE}}(m = 0) = \hat{\mathbb{E}}[Y(1, 0) - Y(0, 0)] = \hat{\alpha}_1$ , is  $-0.32$  under the first approach and  $-0.36$  under the second. Both estimates are statistically significant at the level of 0.01. The estimated CDE if respondents had moral qualms about military intervention, i.e.,  $\widehat{\text{CDE}}(m = 1) = \hat{\mathbb{E}}[Y(1, 1) - Y(0, 1)] = \hat{\alpha}_1 + \hat{\alpha}_3$ , is  $-0.36$  under the first approach but only  $-0.22$  under the second.

### 5.3 Time-Varying Treatment: The Cumulative Effect of Negative Advertising on Vote Shares

We conclude this section with an example of RBW applied to a context involving time-varying treatments and confounders. Political scientists have shown interest in examining the cumulative effect of negative campaign advertising on election outcomes (Lau, Sigelman, and Rovner 2007; Blackwell 2013; Imai and Ratkovic 2015). This is an intricate process because while polling results affect current campaign strategies, they are also constantly shifting, as they respond both to previous results and candidates' use of negative advertising in the past. For their ability to accommodate dynamic causal relationships — particularly since they allow past treatments to affect current outcomes (i.e., “carryover effects”) and past outcomes to influence current treatment (i.e., “feedback effects”) (Imai and Kim 2019) — MSMs are suitable for this research question.

Let  $A_t$  represent our continuous treatment, the proportion of campaign advertisements mentioning the adversary in each campaign week,  $L_t$  our time-varying confounders, the Democratic share and the share of undecided voters in the polls, and  $Y$  the outcome, the Democratic share of the two-party vote. Additionally,

Table 2: MSM Results for the Controlled Direct Effects of Shared Democracy on Public Support for War

	Baseline Confounders Balanced	Baseline Confounders Adjusted in the MSM
Shared democracy	-0.317*** (0.098)	-0.360*** (0.080)
Moral concerns	-1.272*** (0.175)	-1.201*** (0.135)
Shared democracy * Moral concerns	-0.048 (0.210)	0.138 (0.157)

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01

Robust standard errors in parentheses.

we have a set of baseline confounders  $X$  including total campaign length, election year, and whether the election is senatorial or gubernatorial.

Zhou and Wodtke (2020) define an MSM as follows:

$$\mathbb{E}[Y(\bar{a})|X] = \beta_0 + \beta_1 \text{avg}(\bar{a}) + \beta_2 V + \beta_3 \text{avg}(\bar{a})V + \beta_4^T X, \quad (17)$$

where  $V$  is an indicator of incumbency status used to construct interactions that allow the effect to differ between incumbents and nonincumbents, and  $\text{avg}(\bar{a})$  is the average proportion of advertisements that were negative over the final five weeks of the campaign multiplied by ten (we multiply by ten following Zhou and Wodtke (2020) so that regression coefficients can be interpreted as the effect of a ten-percentage point increase in negative advertising).

`rbw` contains two data sets associated with this problem: `campaign_long` and `campaign_wide`. They represent, respectively, the long-format and wide-format data on negative campaign advertising.

The following variables are included in `campaign_long`:

- `neg.dem` is the proportion of advertisements that were negative in a campaign-week.
- `dem.polls` is the Democratic share in the polls.
- `undother` is the share of undecided voters in the polls.
- `neg.dem.l1` is the proportion of advertisements that were negative in the previous campaign-week.
- `dem.polls.l1` is the Democratic share in the polls in the previous campaign-week.
- `undother.l1` is the share of undecided voters in the polls in the previous campaign-week.
- `camp.length` is the length of the candidate's campaign (in weeks).
- `deminc` indicates whether the candidate was an incumbent.
- `base.poll` is the Democratic share in the baseline polls.
- `base.und` is the share of undecided voters in the baseline polls.
- `office` is the type of office in contest:
  - 0: governor
  - 1: senator
- `week` indicates week in the campaign (in the final five weeks preceding the election).
- `year` is the election year.
- `state` is the election state.
- `id` is the candidate id.
- `demName` is the name of the Democratic candidate.

- `demprcnt` is the Democratic share of the two-party vote in the election.

In addition, `campaign_wide` contains the following variables:

- `ave_neg` is the average proportion of advertisements that were negative over the final five weeks of the campaign multiplied by ten.
- `neg.dem_1` is the proportion of advertisements that were negative in week 1 polls.
- `neg.dem_2` is the proportion of advertisements that were negative in week 2 polls.
- `neg.dem_3` is the proportion of advertisements that were negative in week 3 polls.
- `neg.dem_4` is the proportion of advertisements that were negative in week 4 polls.
- `neg.dem_5` is the proportion of advertisements that were negative in week 5 polls.
- `dem.polls_1` is the Democratic share in week 1 polls.
- `dem.polls_2` is the Democratic share in week 2 polls.
- `dem.polls_3` is the Democratic share in week 3 polls.
- `dem.polls_4` is the Democratic share in week 4 polls.
- `dem.polls_5` is the Democratic share in week 5 polls.
- `undother_1` is the share of undecided voters in week 1 polls.
- `undother_2` is the share of undecided voters in week 2 polls.
- `undother_3` is the share of undecided voters in week 3 polls.
- `undother_4` is the share of undecided voters in week 4 polls.
- `undother_5` is the share of undecided voters in week 5 polls.
- `camp.length`, `deminc`, `base.poll`, `base.und`, `office`, `year`, `state`, `id`, `demName`, and `demprcnt` are defined as in `campaign_long`.

Recall from Section 2.3 that RBW requires us to fit a model for the conditional mean of each covariate at each time point given the observed past. We thus estimate regression models of our time-varying confounders  $L_t$  ( $t \geq 2$ ) on lagged treatment and lagged confounders. We also interact each regressor with the week dummies, thus allowing the coefficients to change over time in a flexible manner:

```
x1 <- lm(dem.polls ~ (neg.dem.l1 + dem.polls.l1 + undother.l1) * factor(week),
        data = campaign_long)
x2 <- lm(undother ~ (neg.dem.l1 + dem.polls.l1 + undother.l1) * factor(week),
        data = campaign_long)
```

We then create a list with the model objects to be passed later to the `xmodels` argument in `rbwPanel()`:

```
xmodels <- list(x1, x2)
```

To construct the residual balancing weights, `rbwPanel()` first extracts the residual terms  $\hat{\delta}(L_t)$  from the models above. Note that for each covariate in  $L_1$  (the first time period), the residuals are computed as deviations from the sample mean. Then the function finds a set of weights, such that, in the weighted sample, the residuals are orthogonal to current and future treatments as well as the regressors of  $L_{jt}$ , and the relative entropy between the residual balancing weights and the base weights is minimized. Note that we set the `future` argument to zero to replicate the results from Zhou and Wodtke (2020) since the authors balance the residualized time-varying confounders only with respect to the current treatment, thus assuming away lagged effects of the covariates on the treatment.<sup>5</sup> Since our data do not include sampling weights, a vector of ones is used as the base weights. Additionally, we need to pass arguments indicating the unit id and the time variables due to the longitudinal structure of our data set.

<sup>5</sup>The negative advertising data set spans five weeks, so users interested in considering lagged effects of the time-varying covariates on the treatment may set `future > 0`, up to `future = 4`.

```
rbwPanel_fit <- rbwPanel(treatment = neg.dem, xmodels = xmodels,
  id = id, time = week, data = campaign_long, future = 0)
```

```
## Entropy minimization converged within tolerance level
```

We now attach the weights to `campaign_wide` using the pipe operator and the `left_join()` function from the package `dplyr` (this merging is permitted because the `weights` object returned by `rbwPanel()` is a data frame containing the `id` variable and the residual balancing weights):

```
library("dplyr")
campaign_wide <- campaign_wide %>%
  left_join(rbwPanel_fit$weights, by = "id")
```

We use the functions from `dplyr` to rename the weights so that our variable's name is consistent with the previous examples:

```
campaign_wide <- campaign_wide %>%
  rename(rbwPanel_weights = rbw)
```

Again, we use `survey::svydesign()` to compute robust standard errors:

```
rbwPanel_design <- svydesign(ids = ~1, weights = ~rbwPanel_weights,
  data = campaign_wide)
```

Finally, we use the residual balancing weights to fit the MSM from Equation 17:

```
rbwPanel_msm <- svyglm(demprcnt ~ ave_neg * deminc + camp.length +
  factor(year) + office, design = rbwPanel_design)
```

We report the model results in Table 3. The estimate for nonincumbents is 0.49 and for incumbents is  $0.49 - 1.48 = -0.99$ . The effect of negative advertisement for nonincumbents is positive though not statistically significant. However, the interaction term is significant at the level of 0.01, and incumbents see a sizeable negative effect from negative advertising: a ten percentage point increase in negative advertising reduces a candidate's vote share by about one percentage point, on average.

Table 3: MSM Results for the Cumulative Effect of Negative Advertising on Vote Shares

Average proportion	0.490 (0.315)
Incumbency	14.971*** (2.823)
Average proportion * Incumbency	-1.484*** (0.531)

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01

Robust standard errors in parentheses.

## 6 Conclusion

Compared to other methods of constructing weights for MSMs, RBW has several advantages. In particular, it does not require modeling the conditional distributions of exposures and is thus easy to use with continuous treatments. Previous simulation studies suggest that it is often more efficient and more robust to model misspecification than alternative weighting strategies (Zhou and Wodtke 2020). RBW is also favorable from a computational perspective. Its procedure for finding weights involves a convex optimization problem, allowing RBW to locate a solution substantially faster than alternative methods whose optimization relies on nonconvex loss functions — such as the recently proposed nonparametric version of CBPS, which uses the empirical likelihood (Fong, Hazlett, and Imai 2018). Table 4 sums up these comparisons.

Table 4: Comparison of Methods and Software Implementation

Method	Models for Conditional Distributions of Exposures	Balancing Constraints	Implemented for Time-varying Treatments?	R Package
IPW	Required	Absent	Yes	ipw
CBPS	Required	Present	Yes	CBPS
npCBPS	Not Required	Present	No	CBPS
RBW	Not Required	Present	Yes	rbw

After explaining the underlying logic of RBW, we have described its implementation in the R package `rbw`. With examples from several data sets, we have demonstrated the use of `rbw` in three distinct contexts: effect estimation for point treatments, causal mediation analysis, and effect estimation for time-varying treatments with time-varying confounders.

Nonetheless, RBW is not without limitations. In particular, it depends on models for the conditional means of post-treatment confounders. When these models are incorrectly specified, the pseudo-population generated by residual balancing weights will fail to mimic the original unweighted population, and our estimates will be biased. Even when these models are correctly specified, RBW may also yield biased estimates when we have insufficient balancing conditions. Adding more functions such as cross-products and high-order terms to the set  $G(L_t) = \{g_1(L_t), \dots, g_{J_t}(L_t)\}$ , thereby increasing the number of balancing constraints, may help — though at the risk of making exact balance infeasible. A possible extension to RBW would be to allow for approximate balance with a penalty for the remaining imbalance in the optimization problem (Fong, Hazlett, and Imai 2018).

Despite these limitations, RBW will be useful to many social scientists interested in using marginal structural models to study causality in dynamic settings. We will continue to update the package with functions that expand RBW’s range of applicability (e.g., to contexts involving repeated outcome measures and to cases where, as discussed above, exact balance is infeasible and approximate balance must be pursued).

## References

- Blackwell, Matthew. 2013. “A Framework for Dynamic Causal Inference in Political Science.” *American Journal of Political Science* 57 (2): 504–20. <https://doi.org/10.1111/j.1540-5907.2012.00626.x>.
- Cole, Stephen R., and Miguel A. Hernán. 2008. “Constructing Inverse Probability Weights for Marginal Structural Models.” *American Journal of Epidemiology* 168 (6): 656–64. <https://doi.org/10.1093/aje/kwn164>.
- Fong, Christian, Chad Hazlett, and Kosuke Imai. 2018. “Covariate Balancing Propensity Score for a Continuous Treatment: Application to the Efficacy of Political Advertisements.” *The Annals of Applied Statistics* 12 (1). <https://doi.org/10.1214/17-AOAS1101>.
- Fong, Christian, Marc Ratkovic, and Kosuke Imai. 2021. *CBPS: Covariate Balancing Propensity Score*. <https://CRAN.R-project.org/package=CBPS>.
- Geskus, Ronald B., and Willem M. van der Wal. 2015. *Ipw: Estimate Inverse Probability Weights*. <https://CRAN.R-project.org/package=ipw>.
- Hainmueller, Jens. 2012. “Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies.” *Political Analysis* 20 (1): 25–46. <https://doi.org/10.1093/pan/mpr025>.
- . 2014. *Ebal: Entropy Reweighting to Create Balanced Samples*. <https://CRAN.R-project.org/package=ebal>.
- Imai, Kosuke, Luke Keele, Dustin Tingley, and Teppei Yamamoto. 2011. “Unpacking the Black Box of Causality: Learning About Causal Mechanisms from Experimental and Observational Studies.” *The American Political Science Review* 105 (4): 765–89. <https://doi.org/10.1017/S0003055411000414>.
- Imai, Kosuke, Luke Keele, and Teppei Yamamoto. 2010. “Identification, Inference and Sensitivity Analysis for Causal Mediation Effects.” *Statistical Science* 25 (1): 51–71. <https://doi.org/10.1214/10-STS321>.
- Imai, Kosuke, and In Song Kim. 2019. “When Should We Use Unit Fixed Effects Regression Models for Causal Inference with Longitudinal Data?” *American Journal of Political Science* 63 (2): 467–90. <https://doi.org/10.1111/ajps.12417>.
- Imai, Kosuke, and Marc Ratkovic. 2014. “Covariate Balancing Propensity Score.” *Journal of the Royal Statistical Society. Series B, Statistical Methodology* 76 (1): 243–63. <https://doi.org/10.1111/rssb.12027>.
- . 2015. “Robust Estimation of Inverse Probability Weights for Marginal Structural Models.” *Journal of the American Statistical Association* 110 (511): 1013–23. <https://doi.org/10.1080/01621459.2014.956872>.
- Lau, Richard R., Lee Sigelman, and Ivy Brown Rovner. 2007. “The Effects of Negative Political Campaigns: A Meta-Analytic Reassessment.” *The Journal of Politics* 69 (4): 1176–1209. <https://doi.org/10.1111/j.1468-2508.2007.00618.x>.
- Naimi, Ashley I., Erica EM Moodie, Nathalie Auger, and Jay S. Kaufman. 2014. “Constructing Inverse Probability Weights for Continuous Exposures: A Comparison of Methods.” *Epidemiology (Cambridge, Mass.)* 25 (2): 292–99. <https://doi.org/10.1097/EDE.000000000000053>.
- Robins, James M. 1986. “A New Approach to Causal Inference in Mortality Studies with a Sustained Exposure Period to Control of the Healthy Worker Survivor Effect.” *Mathematical Modelling* 7 (9): 1393–1512. [https://doi.org/10.1016/0270-0255\(86\)90088-6](https://doi.org/10.1016/0270-0255(86)90088-6).
- . 2000. “Marginal Structural Models Versus Structural Nested Models as Tools for Causal Inference.” In *Statistical Models in Epidemiology, the Environment, and Clinical Trials*, edited by Elizabeth M. Halloran and Donald Berry, 95–133. New York: Springer.
- Robins, James M., Miguel A. Hernan, and Babette Brumback. 2000. “Marginal Structural Models and Causal Inference in Epidemiology.” *Epidemiology* 11 (5): 550–60. <https://doi.org/10.1097/00001648-200009000-00011>.

- Tomz, Michael R., and Jessica L. P. Weeks. 2013. "Public Opinion and the Democratic Peace." *The American Political Science Review* 107 (4): 849–65. <https://doi.org/10.1017/S0003055413000488>.
- Urban, Carly, and Sarah Niebler. 2014. "Dollars on the Sidewalk: Should U.S. Presidential Candidates Advertise in Uncontested States?" *American Journal of Political Science* 58 (2): 322–36. <https://doi.org/10.1111/ajps.12073>.
- van der Wal, Willem M., and Ronald B. Geskus. 2011. "Ipw: An R Package for Inverse Probability Weighting." *Journal of Statistical Software* 43 (13): 1–23.
- Vansteelandt, Stijn. 2009. "Estimating Direct Effects in Cohort and CaseControl Studies." *Epidemiology (Cambridge, Mass.)* 20 (6): 851–60. <https://doi.org/10.1097/EDE.0b013e3181b6f4c9>.
- Wang, Yue, Maya Petersen, David Bangsberg, and Mark van der Laan. 2006. "Diagnosing Bias in the Inverse Probability of Treatment Weighted Estimator Resulting from Violation of Experimental Treatment Assignment." *U.C. Berkeley Division of Biostatistics Working Paper Series Working Paper 211* (September).
- Zhou, Xiang, and Geoffrey T. Wodtke. 2020. "Residual Balancing: A Method of Constructing Weights for Marginal Structural Models." *Political Analysis* 28 (4): 487–506. <https://doi.org/10.1017/pan.2020.2>.